

Supplementary Information for “Matter that learns what matters”

N. G. Berloff

Department of Applied Mathematics and Theoretical Physics, University of Cambridge
Cambridge CB3 0WA, United Kingdom

Guide to the Supplementary Information

This Supplement presents the model, derivations and numerical tests supporting the article. Section S1 specifies the dynamics and decision protocol. Sections S2–S6 derive the analytical results. Section S7 reports the numerical tests, grouped by the claims they examine. Section S8 sets out the requirements for a physical implementation. Section S9 gives the accuracy checks and lists the data and code to be deposited. The historical test sets R1–R5 (criteria fixed before the runs; 62 registered statements) are reported only where the article uses them; the complete registry, including the statements that failed, and the material removed from this Supplement during the revisions will be deposited with the data and code (Section S9). Later studies state their criteria in their own subsections. Equation numbers prefixed by S refer to this Supplement; unprefixed numbers refer to the article, whose Methods equations begin at (5).

Where to find the evidence for each article statement

Article statement	Supplement
Element, bistability and its folds (article Fig. 1a)	Section S1.1
Two readouts of one record: sensing signal and reader gain (article Fig. 1a)	Sections S1.3, S3.1
Network, decision cycle, parameters, noise	Sections S1.3, S1.5; Table S2
What the traces credit; why use credits confirming readings less	Section S2
Sensing bound, summed-gain protection, valid revision, width range	Section S3.3
Functional window at $w = 0.6$ for valid records in [3.30, 3.70]	Section S3.4
Retention while a need is inactive; assignment symmetry	Sections S4.1, S4.2
Credit criterion and learning-retention target	Sections S5.2, S5.3
Self-limiting writer: boundary and design conditions	Section S6; Fig. S3
Design, statistics and criteria	Section S7.1
Formation, false credit, co-reading (article Fig. 3a–c)	Section S7.2
Transplant (article Fig. 2c,d)	Section S7.3
One network learning in time (article Fig. 2a,b)	Section S9.2; Fig. S2
Retention, revision, durability, reader withdrawal (article Fig. 3a,b,d,e)	Section S7.4
Robustness to parameters, noise, code and network size	Section S7.5
Frozen-record tests (article Fig. 4b,c)	Section S7.6
Learning misses the sensing interval from both sides (FPL1; article Fig. 4d)	Section S7.7
A correct outcome pushes a record out (FRCP1; article Fig. 4e)	Section S7.8
Sensing lifetimes under continued learning	Section S7.9
Self-limiting writer in the network (FML1)	Section S7.10
Six-network validation, capacity effect and the limit on replacement (S3VAL; article Fig. 5a–d)	Section S7.11
Reset of the obsolete record; adverse-outcome susceptibility (RLR1)	Section S7.12
Fresh three-arm comparison of the susceptibility, the self-limiting writer and a uniform control (VAL1; article Fig. 5e)	Section S7.13
Candidate polariton link, experimental precedents and missing steps, mapping to the equations, decisive tests	Sections S8.1–S8.3
Step-size checks including the step-halving check (RES1), markers † and ‡, data and code	Section S9

Notation

A *need* is a yes-or-no question about one hidden binary value of the environment. The two needs, E and I, occupy the two rows of article Fig. 1. Each row connects the three sources to one need’s answer memory. A source is *valid* for a need when it reports that need’s current hidden value. An *assignment* specifies which source is valid for each need. In assignment A, source 1 is valid for E, source 2 is valid for I and source 3 is unrelated to both. The network is not told the assignment. Symbols retain their meanings from the article. Table S1 lists those used most often here, together with those introduced in this Supplement.

Study labels identify archived test sets; network labels identify individual networks within a study. Archive and registered-statement identifiers retain their original numbering, so selected summaries need not contain every number. Unless stated otherwise, the cavity detuning is $\delta = 3$ and the holding power is $A_0 = 4$.

Table S1: Notation used in this Supplement.

Term	Symbol	Meaning
Branches	x_L, x_U, x_H	Low, unstable and high quiet populations; $x_L = 2 - \sqrt{2}$, $x_U = 2$, $x_H = 2 + \sqrt{2}$. A record is formed when $x > x_U$.
Formed pattern; validity	$S_{jk}; V_{jk}$	$S_{jk} = 1$ when $x_{jk} > x_U$. $V_{jk} = 1$ when source k is valid for need j ; V is used only for evaluation.
Functional window	$\mathcal{G}_V$	Record states in which sensing, protection and revision all hold for assignment V ; protection and revision refer to ordinary (unassisted) readings.
Readout targets	$p_*; m; q_{\max}$	Minimum valid-source selection probability; memory decision margin; outer bound on $ \text{Re } z_j $.
Sensing limits; folds	$x_-, x_+; x_{f-}, x_{f+}$	Lower and upper ends of the sensing interval (Section S3.3); turning points of the stationary response $A = xD(x)$.
Summed wrong gain; residual	$\overline{G}_d; \eta_r$	Bound on the sum of active non-valid-reader gains; bound on the terms omitted from the reduced memory equation.
Source value; its sign; hidden value	$y_k = 3\zeta_k; \zeta_k; \zeta_j$	y_k is the field of source k during a reading and $\zeta_k = \pm 1$ its binary sign; $\zeta_j = \pm 1$ is the current hidden value of need j , which its valid source reports.
Writer coordinate; trace area	$\Psi; \mathcal{A}$	$\Psi = x + x_* \ln(x/x_*)$ makes the writer additive; $\mathcal{A} = \Lambda \int \sigma \epsilon dt$. $\mathcal{A}$ is not the drive power A .
Outcome threshold	ν	$\nu = \beta_{j*} + \vartheta$.
Write line; capacity	$a_w(x); \kappa_{\text{cap}}(x); x_a, x_c$	Factors of the self-limiting writer; their product is the positive factor f_w . The line with the capacity (PC) is the self-limiting writer of article equation (4). κ_{cap} is not the cavity decay rate κ .
Positive-branch coordinate	Φ	$\Phi = \int dx / (H a_w \kappa_{\text{cap}})$; additive for positive pulses of the self-limiting writer; diverges at x_c .
Memory coordinates; potential	$q_j, p_j; \mathcal{V}$	$z_j = q_j + ip_j$; $\mathcal{V}$ is the memory potential, distinct from the reservoir potential U_A .
Time units	$t; \tau = \gamma t$	t in memory-response units. A decision cycle lasts 22 units; a need-cycle (one action of a given need) lasts 44.
Action; score	s	The answer given by the acting need, scored $s = 1$ if correct and 0 otherwise.
Writer; credit	Π_{jk}	The outcome-dependent update of the records, article equation (3); a link’s credit is its share of an outcome’s write.
Arm		One writer law or condition within a study.
Panel; assay		A set of readout tests run on copies of the network state at a given cycle; an assay is one such test.
Network		One seeded simulation of a network under one condition (called a family in the labels of studies, for example F02). Links within a network are not independent.
Evidence sets	R1–R5; G1–G3	Historical registered test sets; G1 replays and rescores them; G2 is a fresh matched cohort; G3 continues G2 with interventions.
Adverse-outcome susceptibility	$\mu_{jk}; r_\mu, T_\mu, \eta_\mu$	Second state of a link in the extension of Section S7.12, Eq. (S7.1); it is not the memory margin m .

Term	Symbol	Meaning
Studies at $w = 0.6$	FPL1; FRCP1; FML1; S3VAL; RLR1; VAL1	Learning into the functional window; matched outcome-pulse replays; self-limiting writer study in a re-implementation; its six-network check in the original code (S3VAL networks N1–N6); the reset and the adverse-outcome extension on those networks; the fresh three-arm comparison. Histories are labelled by trace rule and family, for example U_F02 (use, family 2).
Trace rules	C, A, U	Co-presence, admission, use. Markers $\dagger$ and $\ddagger$ are defined in Section S9.1.
Durability	O_f	Correct-only cycles in 2501–3000; a durable run has $O_f \geq 400$ and no persistent false record.
Registered; post-hoc		Criteria fixed before the runs; analyses chosen after inspecting results.

Contents

Notation	2
S1 Model, parameters and noise	4
S1.1 The element	4
S1.2 Probe and read tones	6
S1.3 Network and decision cycle	6
S1.4 Parameters	8
S1.5 Noise	9
S2 What the three traces credit	9
S2.1 Co-presence and admission	10
S2.2 Use is effective work done on the memory	10
S2.3 Credit generated by the reading	11
S3 Readout bounds and the functional window	11
S3.1 Sensing and reader gain at the two branches	11
S3.2 Why learning needs assisted reading	12
S3.3 Comparison conditions for simultaneous readings	12
S3.4 The full-reader argument	14
S4 Retention and separation	16
S4.1 A finite residual-credit bound for dormancy	16
S4.2 Assignment symmetry at fixed network size	17
S5 Finite-pulse credit and the learning-retention target	17
S5.1 The outcome baseline and the margin	17
S5.2 Exact trace-weighted credit for finite pulses	18
S5.3 The learning-retention target	19
S6 The self-limiting writer	19
S7 Simulation evidence	22
S7.1 Design, statistics and criteria	23
S7.2 Learning and false credit	23
S7.3 Transplant	24
S7.4 Retention, revision and continued use	24
S7.5 Robustness	25
S7.6 The functional window on stored states	26
S7.7 Learning misses the sensing interval from both sides (FPL1)	26
S7.8 A correct outcome causes the exit (FRCP1)	27
S7.9 Sensing lifetimes under continued learning	27
S7.10 The self-limiting writer in the network (FML1)	28
S7.11 Six-network validation in the original simulation code (S3VAL)	30
S7.12 Releasing the obsolete record (RLR1)	32
S7.13 Fresh three-arm comparison (VAL1)	34
S8 Physical implementation	35
S8.1 A candidate polariton learning link and its experimental support	35
S8.2 Mapping to the equations	37
S8.3 Decisive experimental tests	38
S9 Numerical accuracy, data and code	39
S9.1 Integration step and the result markers	39
S9.2 Replay and independent implementations	40
S9.3 Data and code	40

S1 Model, parameters and noise

This section explains how a link stores a relevance record, how that record is probed without being rewritten and how the links form the network. It also specifies the decision protocol, parameters and noise model. These details support the article’s model description, Methods equations (5)–(11) and article Fig. 1a.

S1.1 The element

Each network link contains an optical cavity coupled to a reservoir of longer-lived excitations. The reservoir population n stores the link’s relevance record. The binary answer is stored separately, in the need’s answer memory (Section S1.3). Absorbed cavity light fills the reservoir, whose excitations shift the cavity resonance. This shift changes subsequent absorption and allows two stable populations under the same drive. Reference [1] derives coupled condensate–reservoir equations for non-resonantly pumped polariton condensates. It gives general forms of the reservoir term and the regimes obtained by varying the reservoir relaxation rate relative to that of the condensate. Here we consider a resonantly driven mode in the slow-reservoir limit. Eliminating the fast field leaves the reservoir population as the dynamical variable. The holding drive remains on while the record is retained, so driven feedback maintains the record.

In a frame rotating with the drive, let a be the complex cavity amplitude and f the drive. The drive frequency lies $\Delta > 0$ above the empty-cavity resonance. The reservoir raises this resonance by $g_R n$, with $g_R > 0$, so an increasing population initially brings the cavity towards the drive frequency. The cavity amplitude decays at rate κ . Capture at rate s_{cap} (distinct from the action score s) fills the reservoir at rate $s_{\text{cap}}|a|^2$. Reservoir excitations are lost at rate γn :

$$\dot{a} = [-\kappa + i(\Delta - g_R n)]a + f, \quad \dot{n} = s_{\text{cap}}|a|^2 - \gamma n. \quad (\text{S1.1})$$

Capture is part of the cavity loss. With unit capture yield and $|a|^2$, n counted in excitations, each capture removes one cavity excitation and adds one reservoir excitation. Because κ is an amplitude-decay rate, the undriven cavity population decays at rate 2κ , so $\kappa = \kappa_{\text{other}} + s_{\text{cap}}/2$ with $\kappa_{\text{other}} \geq 0$, and $s_{\text{cap}} \leq 2\kappa$.

We measure the reservoir-induced shift in units of κ and rescale the field and the drive so that the reservoir filling rate becomes $\gamma|b|^2$:

$$x = \frac{g_R n}{\kappa}, \quad b = \sqrt{\frac{g_R s_{\text{cap}}}{\gamma \kappa}} a, \quad \delta = \frac{\Delta}{\kappa}, \quad \tilde{f} = \sqrt{\frac{g_R s_{\text{cap}}}{\gamma \kappa^3}} f, \quad A = |\tilde{f}|^2, \quad D(x) = 1 + (\delta - x)^2. \quad (\text{S1.2})$$

Then

$$\dot{b} = \kappa[(-1 + i(\delta - x))b + \tilde{f}], \quad \dot{x} = \gamma(|b|^2 - x). \quad (\text{S1.3})$$

Here $\delta - x$ is the detuning between the drive and the shifted cavity resonance. The factor $D(x)$ is smallest at resonance, $x = \delta$. A smaller D gives a larger stationary optical response at the same drive power. In a stationary state of Eq. (S1.3) under constant drive, filling balances reservoir loss and $|b|^2 = x$. During a transient, these quantities can differ substantially (Fig. S1c). We call the element *quiet* when only the holding drive A_0 acts, without weak tones, reader light or writing. Its stationary populations are the quiet populations defined below.

Reduction and equilibria. When $\kappa \gg \gamma$, the cavity field settles much faster than the reservoir changes. After the initial field transient its intensity therefore follows the current population, $|b|^2 \simeq A/D(x)$, and the reservoir obeys one equation,

$$\frac{\dot{x}}{\gamma} = F_A(x) = -x + \frac{A}{D(x)} = -U'_A(x), \quad U_A(x) = \frac{x^2}{2} + A \arctan(\delta - x). \quad (\text{S1.4})$$

This is article equation (1). At the holding power we write $F_4 = F_{A_0}$ with $A_0 = 4$. The two terms of F_A describe reservoir loss and optical replenishment. The minima of the effective potential U_A give the stable populations. At a stationary state, loss balances replenishment, giving $A = xD(x)$. This relation is exact for the full element: it also follows by setting both derivatives in Eq. (S1.3) to zero. The separation of time scales is needed only to describe the transient with one equation.

For positive detuning the stationary response $A = xD(x)$, a cubic in x , has two turning points when $dA/dx = 3x^2 - 4\delta x + \delta^2 + 1 = 0$ has real roots, that is for $\delta > \sqrt{3}$. At each turning point a stable branch meets the unstable branch and disappears. At $\delta = 3$ the turning points and their powers are

$$x_{f\pm} = \frac{1}{3}(2\delta \pm \sqrt{\delta^2 - 3}) = 2 \pm \sqrt{\frac{2}{3}} = 1.183503, 2.816497, \quad A(x_{f-}) = 5.088662, \quad A(x_{f+}) = 2.911338. \quad (\text{S1.5})$$

The reduced element is therefore bistable for drive powers between 2.911 and 5.089. The same drive can hold either population, depending on the history. At the holding power $A_0 = 4$, the cubic factorises exactly as $(x - 2)(x^2 - 4x + 2) = 0$, giving

$$x_L = 2 - \sqrt{2} \simeq 0.5858, \quad x_U = 2, \quad x_H = 2 + \sqrt{2} \simeq 3.4142. \quad (\text{S1.6})$$

The slopes F'_4 at these points are $-(2 - \sqrt{2})$, $+1$ and $-(2 + \sqrt{2})$, so small disturbances decay at the outer states and grow at the middle one. In the scalar dynamics at the holding power, x_U separates the populations that return to x_L from those that return to x_H . A pulse that carries the population across x_U therefore leaves a record after the pulse ends, while the holding drive stays on. In the network, $x > x_U$ is the criterion that classifies a record as formed; a formed record need not have useful readouts (Section S3.3).

The element in time. Figure S1 shows solutions of the full field and reservoir equations (S1.3) with $\kappa = 100$ and $\gamma = 0.001$ in memory-response units. Panel (a) shows switching through changes of the holding power. A pulse to $A = 8$ exceeds the upper fold at 5.089, where the low branch ends, and carries the population across x_U . Returning to $A_0 = 4$ then retains the high branch. A dip to $A = 3.3$, above the lower fold at 2.911, lowers the record to about 3.2 without erasing it. A pulse to $A = 2$ removes the high branch and carries the record back below x_U . These changes of holding power illustrate the element's behaviour. During learning, records instead change through the outcome writer of article equation (3).

Panels (b) and (c) show that the optical field responds much faster than the record. With both weak tones on for ten time units (panel b), the cavity population rises by about 0.68 at a formed record and 0.12 at an unformed one. Each reservoir population changes by less than 0.01, and both records stay on their branches. The probe thus distinguishes the stored states without rewriting them. Panel (c) resolves the cavity transient at the start of the write pulse, during which the record is essentially unchanged.




Figure S1: **One element in time.** Solutions of the full equations (S1.3). (a) Switching through changes of holding power. The record x (blue), cavity population $|b|^2$ (red) and drive power A (grey dashed) are shown through a pulse that writes the record, a dip that leaves it on the high branch and a pulse that erases it. Dotted lines mark x_L , x_U , x_H ; time is in reservoir lifetimes γt . (b) Both weak tones act for one ten-unit reading interval (shaded) at an unformed (blue) and a formed (red) record. The cavity populations $|b|^2$ (solid) distinguish the records while the reservoir populations x (dotted) change only slightly. (c) Fast cavity response to the write pulse beginning at $\kappa t = 20$. Time κt counts field-amplitude relaxation times; the record x is essentially unchanged on this scale. The curves are deterministic, with no error bars.

Formation and erasure face different restoring forces. For a constant additional reservoir input $\pi = \Pi/\gamma$ in the reduced equation, formation from x_L requires π to exceed the largest opposing drift on $[x_L, x_U]$, $\max(-F_4) = 0.2764$, whereas erasure from x_H requires $\pi < -\max_{[x_U, x_H]} F_4 = -1.0635$, about 3.85 times larger in magnitude. A writer must therefore overcome a much stronger flow to erase a formed record than to form one. The finite, trace-dependent outcome pulses of learning are analysed in Section S5; the constant-input thresholds do not apply to them.

Local stability of the full element. The two storage states remain stable when the cavity field evolves dynamically. At $A_0 = 4$ and $\delta = 3$, the Routh–Hurwitz conditions for the three-dimensional linearisation of Eq. (S1.3) hold at x_L and x_H for every $\kappa > 0$ and $\gamma > 0$. The middle state x_U has one unstable direction. Local stability therefore does not require a large separation of time scales. Section S1.5 treats noise-induced transitions.

S1.2 Probe and read tones

Two weak tones are added to the holding drive of every link: the probe (priority pump) on the row of the acting need and the read tone on the column of the source being read. They change the optical response without switching the stored branch.

Quadrature addition. For one link the drive is $\tilde{f} = \sqrt{A_0} + i\sqrt{B}(P + I)$ with $P, I \in \{0, 1\}$. The tones are in phase with each other and in quadrature with the holding drive. Then

$$|\tilde{f}|^2 = A_0 + B(P + I)^2 : \quad 4.0, 4.2, 4.8 \quad \text{for } P + I = 0, 1, 2 \quad (B = 0.2). \quad (\text{S1.7})$$

Tones in phase with the holding drive would give instead $A_0 + 2\sqrt{A_0 B}(P + I) + B(P + I)^2 = 5.988854$ and 8.377709 . One tone alone would then exceed the upper fold 5.088662 , and two active tones, over repeated readings of a frequently selected source, would carry its record across the fold without any outcome. The quadrature is therefore part of the element.

S1.3 Network and decision cycle

Needs have indices $j \in \{E, I\}$, as in article Fig. 1. Sources have indices $k \in \{1, 2, 3\}$, and j^* is the acting need. The deterministic equations are

$$\dot{b}_{jk} = \kappa\{-1 + i(\delta - x_{jk})\}b_{jk} + \tilde{f}_{jk}, \quad \tilde{f}_{jk} = \sqrt{A_0} + i\sqrt{B}(P_j + I_k), \quad (\text{S1.8})$$

$$\dot{x}_{jk} = \gamma(|b_{jk}|^2 - x_{jk} + \chi|c_{jk}|^2) + \Pi_{jk}(t), \quad (\text{S1.9})$$

$$\dot{c}_{jk} = -[\kappa_r + i\Delta_{r,jk}(t)]c_{jk} - ivI_k(z_j - y_k), \quad (\text{S1.10})$$

$$\dot{z}_j = \left[\frac{1 - G}{2} - |z_j|^2 \right] z_j + \frac{1 + G}{2} z_j^* - i \sum_k v I_k c_{jk} + u_j(t). \quad (\text{S1.11})$$

Here κ_r is the reader decay rate and v its coupling. The parameter G damps the imaginary memory quadrature, and χ gives the share of reader light that fills the reservoir. The weak tones have power scale B ; P_j controls the probe, I_k the source switch and u_j the correction drive. Table S2 gives their values. Reservoirs have no direct coupling. In an ordinary reading, $\Delta_{r,jk} = \Delta_r(x_{jk}) = \kappa_r(x_{jk} - x_H)/w$ and $g_0 = v^2/\kappa_r$. Assisted reading sets the selected reader of the acting need to zero detuning. This coefficient change is supplied by the protocol. The writer Π_{jk} is article equation (3), driven by one of the traces of Methods equations (12)–(14). Time is measured in memory-response units.

The source value is $y_k = 3\zeta_k$ during a reading and zero otherwise; $\zeta_k \in \{-1, +1\}$ is the binary sign that source k reports, fixed by the assignment. After scoring, the correction drive writes the scored hidden value of the acting need. For a binary action $\mathbf{a}$ and score $s = \mathbf{1}\{\mathbf{a} = \zeta_{j^*}\}$ this value is exactly

$$\zeta_{j^*} = \mathbf{a}(2s - 1), \quad u_{j^*} = 3\mathbf{a}(2s - 1). \quad (\text{S1.12})$$

The correction therefore carries no task information beyond the known action and the binary score. The other memory receives no drive.

Preparation. Each run sets $b = c = z = \epsilon = 0$ and $x = x_L$, then evolves for 100 quiet time units with noise. The answer memories stay at zero during this time because they receive neither noise nor a reader drive. The task then draws the hidden values and sets $\beta_j = 0.5$.

One decision cycle. Times are measured from the start of a cycle of the acting need j^* ; the other probe is off throughout.

Time	Controls	Operation
0–2	$P_{j^*} = 1$; all $I_k = 0$	Sense; sample D_k at $t = 2$ and draw the selected source.
2–12	$P_{j^*} = 1$; selected $I_k = 1$	Read. Assisted reading, when drawn, acts only on (j^*, k) , never on a co-read source.
12	Readers and assistance off	First hidden-value flip, probability 0.1 per value in the flip protocol.
12–14	$P_{j^*} = 1$	Delay with readers disconnected.
14	Score, then $P_{j^*} = 0$	Draw the designated guess, choose the action and score it against the hidden value after the flip. Update the chance baseline where applicable.
14–16	$u_{j^*} = 3a(2s - 1)$	First half of the correction. $\sigma = 1$ after non-guess actions, with the same outcome factor on every link.
16–18	u_{j^*} on; $\sigma = 0$	Second half of the correction.
18–22	All controls off	Quiet evolution.
22	Second hidden-value flip	Alternate the acting need.

Designated guesses occur with probability 0.2 and use a fair coin. They update only $\beta_{j^*} \leftarrow (1 - \alpha)\beta_{j^*} + \alpha s$, with $\alpha = 0.02$, and receive no outcome pulse. A non-guess action uses a fair coin when $|\text{Re } z_{j^*}| < 0.2$; such an action still writes records. Because the probe is off during the pulse, admission and use integrate a decaying trace over 14–16. A decision cycle lasts 22 time units; a need-cycle is one action of a given need and spans 44 units. A copied valid source value faces one flip before scoring and succeeds with probability 0.9. An answer kept from the previous action of the same need faces four flips and agrees with probability $[1 + (1 - 2p_{\text{flip}})^4]/2 = 0.7048$, with flip probability $p_{\text{flip}} = 0.1$. With the 20% designated guesses, a network without readers that always keeps the corrected answer therefore succeeds in about 0.664 of actions.

Protocol variants. The redraw protocol draws new hidden values at $t = 0$ of every cycle and has no mid-cycle flips; a valid source value then predicts the score exactly, and an old answer is correct with probability 1/2. In the co-reading test, whenever the selected source is valid for the acting need, the unrelated source is also read over 2–12 with its own source value. It never receives assisted reading. The null test has no valid source. R1–R4 feed all sources with one shared unrelated value; R5, G2 and G3 use independent unrelated values per source, so their null results are reported separately.

Figure S2 shows three decision cycles of one learning network, the admission network of article Fig. 2a,b: an assisted reading before the record forms, an ordinary reading through the still-unformed record, and an ordinary reading once the record has formed.




Figure S2: **Three decision cycles of one learning network.** The admission network of article Fig. 2a. In each cycle, need E reads its valid source, source 1, with source value +3 and stored answer -1 . Rows show the three need-E cavity populations $|b_{Ek}|^2$ and the probability of selecting source 1; the answer memory $\text{Re } z_E$ and reader gain relative to resonance, $g(x)/g_0$; and record change with peak admission trace. Shading marks sensing (0–2), reading (2–12) and the outcome pulse (14–16). The vertical line marks the action. **a**, Cycle 25: assisted reading tunes the reader to resonance while the record is unformed ($x = 0.56$). The answer reverses, the action is correct and the record rises by 0.09. **b**, Cycle 111: an ordinary reading through the unformed record ($x = 1.04$, gain 0.15) moves the memory only part way. The action is wrong, and the same trace lowers the record by 0.18. **c**, Cycle 701: the formed record ($x = 3.96$) raises source 1’s sensing signal. Its selection probability is 0.88, against 0.40 in **b**, and its reader gain is 0.77 without assistance. The answer reverses, the action is correct and the record rises by 0.18. After each action, the correction drive sets the memory to the hidden value. The cycles were selected for their common source value and initial answer sign. Full network equations with noise, recomputed with a passive observer. Reader gain is article equation (2) at the simulated record; selection probability is article equation (11) applied to the simulated sensing signals. One realisation; no error bars.

S1.4 Parameters

Table S2 lists the reference parameters. Deviations from the reference parameters are stated with each study.

Table S2: Reference parameters.

Symbol	Value	Meaning
A_0, δ	4, 3	Holding power and cavity detuning
B	0.20	Power scale of each weak tone
κ, κ_r	100, 200	Cavity and reader amplitude-decay rates
γ	0.001	Reservoir decay rate
g_0, w	0.7, 1.0	Peak reader gain and reader width
χ	0.01	Reader contribution to reservoir filling
G	80	Damping of the imaginary part of the memory
T_ϵ, ϵ_0	10, 0.01	Trace lifetime and cutoff
S_a, S_u	10, 1	Admission and use signal scales
Λ, τ_w	10, 2	Writer gain and pulse duration; $\Lambda\tau_w = 20$
x_*, ϑ, α	1, 0.2, 0.02	Writer response scale, outcome margin and baseline update fraction
β_j (initial)	0.5	Chance baseline
p_{as}	0.5	Probability of assisted reading
β_s, ϵ_x, N_K	8, 0.1, 3	Selection preference, exploration fraction and number of sources
p_{guess}	0.2	Probability of a designated guess
$ y_k , u_{j*} $	3, 3	Source and correction field amplitudes
p_{flip}	0.1	Hidden-value flip probability at $t = 12$ and $t = 22$
$T_{\text{sense}}, T_{\text{read}}$	2, 10	Sensing and reading durations
$T_{\text{delay}}, T_{\text{corr}}, T_{\text{quiet}}$	2, 4, 4	Delay, correction and quiet durations (the outcome pulse lasts $\tau_w = 2$)
D_τ, θ_b	0.01, 10^{-4}	Additive reservoir and mode noise
Δt	0.002	Reference integration step

R3 varies one parameter at a time: $\delta \in \{2.95, 3.05, 3.10\}$, $B \in \{0.15, 0.25\}$, $\Lambda\tau_w \in \{15, 30\}$, $x_* \in \{0.5, 2\}$, $w \in \{0.8, 1.2\}$ and $p_{as} \in \{0.3, 0.6\}$. Together with the reference point, these give 14 settings. R1 and R2 also test $\gamma = 0.0005, 0.002, 0.004$, $p_{as} = 0.2, 0.35$ and $\Lambda\tau_w = 10$. Coupled parameter changes were not tested.

S1.5 Noise

Independent real Wiener increments are added to each field quadrature and each record:

$$db_{jk} = \dot{b}_{jk}^{\text{det}} dt + \sqrt{2\theta_b} (dW_{jk}^{b,1} + i dW_{jk}^{b,2}), \quad dx_{jk} = \dot{x}_{jk}^{\text{det}} dt + \sqrt{2D_\tau\gamma} dW_{jk}^x,$$

with $D_\tau = 0.01$ and $\theta_b = 10^{-4}$; the field-noise amplitude is not multiplied by κ . The increments are applied after each fourth-order Runge–Kutta drift step, and a negative record is set to zero. The readers, memories and traces receive no direct noise. The use signal evaluates the deterministic $\dot{z}_j$ at each Runge–Kutta stage. Since $P_j = 0$ while $u_j \neq 0$, the correction adds no use-trace input.

For the quiet element in reservoir time, $dx = F_4(x) d\tau + \sqrt{2D_\tau} dW_\tau$, additive reservoir noise can carry the population between the two stable states. The rises of the effective potential U_4 of Eq. (S1.4) from each stable population to the unstable one are

$$\Delta U_L = U_4(x_U) - U_4(x_L) = 0.257631, \quad \Delta U_H = U_4(x_U) - U_4(x_H) = 0.883962, \quad (\text{S1.13})$$

so the quiet storage states are metastable under weak noise.

R3 also replaces additive record noise by finite-step Poisson birth and death increments with rates $\Omega\{\gamma(|b_{jk}|^2 + \chi|c_{jk}|^2) + [\Pi_{jk}]_+\}$ and $\Omega\{\gamma x_{jk} + [-\Pi_{jk}]_+\}$. At $\Omega = 50$, about 29 and 171 excitations at the quiet populations, all eight admission runs formed both valid records, no persistent false record appeared among 32 non-valid or 48 null links, and all eight dormant records met the floor 2.5 criterion. At $\Omega = 12$, all eight runs still formed both valid records, but six of 32 non-valid links acquired persistent false records, the null count was one of 48, and five of eight holds met the dormancy criterion (Table S7).

S2 What the three traces credit

The writer can credit only what its trace measures. This section derives the quantity measured by each trace in Methods equations (12)–(14). It explains why use credits a confirming reading about one tenth as much as a reversing one. It also bounds the credit given to an unrelated co-read source relative to the valid source and identifies when the unrelated source still receives credit. A delayed broadcast outcome acts on a transient local trace. This is the architecture of three-factor synaptic plasticity: experiments show delayed dopamine

acting on a state left by recent synaptic activity [2, 3]. Non-equilibrium physical dynamics can likewise retain the transient information needed by a local learning rule [4]. A trace enables delayed local credit but does not establish causal responsibility. Sections S2.2 and S5.2 determine when its credit is informative.

S2.1 Co-presence and admission

The co-presence trace uses the instantaneous cavity intensity, including its noise, before the cutoff. The trace itself is the low-pass filter. At a fixed record with the field settled, one tone contributes $B/D(x)$ and both tones contribute $4B/D(x)$ before the cutoff:

	One tone	Both tones
Low branch x_L	0.029289	0.117157
High branch x_H	0.170711	0.682843

A high link with one tone can therefore have more trace input than a low link with both tones. Co-presence measures light, not whether a reading changed the answer. At $\epsilon_0 = 0.01$ the one-tone to two-tone ratio is 0.180 on the low branch and 0.239 on the high branch.

The admission trace measures reader activity. The reader responds to the field mismatch $z_j - y_k$. Because the source field has amplitude 3 and the quiet memory has amplitude 1, a source value that confirms the stored sign still drives the reader. Admission therefore credits the link that was read, whether or not the reading changed the answer.

All three traces use $T_\epsilon = 10$ and $\epsilon_0 = 0.01$, and one trace rule drives the writer in a run. A ten-unit constant input followed by a two-unit delay leaves a fraction $(1 - e^{-1})e^{-0.2} = 0.52$ of the sustained trace at the pulse.

S2.2 Use is effective work done on the memory

Work in this section is work in the effective memory potential; it is not a measure of optical or thermodynamic energy, which would require the physical reservoirs and entropy flows to be identified [5]. Write $z_j = q_j + ip_j$ and $r_{jk} = -ivI_k c_{jk}$ for the reader's contribution to $\dot{z}_j$. The memory dynamics without inputs is a gradient flow with potential

$$\mathcal{V}(q, p) = \frac{1}{4}(q^2 + p^2)^2 - \frac{1}{2}q^2 + \frac{G}{2}p^2. \quad (\text{S2.1})$$

The real inner product of the memory equation with its velocity gives exactly

$$\sum_k \text{Re}(\dot{z}_j^* r_{jk}) = \frac{d\mathcal{V}(q_j, p_j)}{dt} + |\dot{z}_j|^2 - \text{Re}(\dot{z}_j^* u_j). \quad (\text{S2.2})$$

During reading, $u_j = 0$, giving Methods equation (15). This identity uses the full reader and memory equations. Gating by the row, applying the cutoff and rectifying each term alter the sum. The resulting division of work among readers is therefore approximate. The signal measures alignment with memory motion; it does not assign exact responsibility for success.

Why confirming readings receive less credit. Consider one real memory coordinate q . Eliminate the fast reader field as in Section S3, and hold its real gain g and source value y constant. The reader contribution is then $r = g(y - q)$, so

$$\int_0^T r \dot{q} dt = \frac{g}{2}[(q(0) - y)^2 - (q(T) - y)^2]. \quad (\text{S2.3})$$

At $g = 0.7$, $y = -3$, the stable endpoint is $q_\infty = -1.358577913$, the negative root of $q^3 - 0.3q + 2.1 = 0$. A confirming reading starts at $q(0) = -1$ and a reversing reading at $q(0) = 1$. As $T \rightarrow \infty$ their integrated signals are 0.457006737 and 4.657006737, a ratio 0.098133149. With the trace filter ($T_\epsilon = 10$, $\epsilon_0 = 0.01$, $S_u = 1$, a ten-unit reading and a two-unit delay) the pulse-time traces are 0.01365070 and 0.14650841, a ratio 0.09317347. At a frozen low record, an opposing source value moves the memory from 1 to about 0.79 without changing its sign, and the trace at the pulse is about 0.001. These are scalar-reduction values, not network averages.

A bound for an unrelated co-read source. In the real fixed-gain reduction, let the valid reader have gain g_v and an unrelated reader gain $0 \leq g_d \leq g_v$. Let both source values have magnitude V . Assume that $q \in (-V, V)$ moves monotonically towards the valid value and that both links have the same row gate and use scale. Set $\eta_g = g_d/g_v$. When the source values agree, $r_d = \eta_g r_v$. Since $[\eta_g h - \epsilon_0]_+ \leq \eta_g [h - \epsilon_0]_+$ for $0 \leq \eta_g \leq 1$, filtering the traces gives

$$\epsilon_d^u(t) \leq \eta_g \epsilon_v^u(t) \quad (\text{S2.4})$$

whenever the initial traces satisfy this inequality. If the source values oppose, the unrelated push opposes the motion and adds no use input; an earlier unrelated trace still decays. At $g_v = 0.7$, $g_d = 0.7/9$, agreeing source values give the bound $1/9$. For opposing source values, the scalar memory crosses zero monotonically if $V(g_v - g_d) > \frac{2}{3\sqrt{3}}(1 - g_v - g_d)_+^{3/2}$, with margin 1.82634585 at $V = 3$. The bound concerns relative trace size, not the probability that a false record forms; complex phases, moving records and later updates need not satisfy its premises.

S2.3 Credit generated by the reading

In the network the trace input comes from the same reading that changes the answer. Freeze the records during one reading, eliminate the fast readers and keep the real memory coordinate $q = \text{Re } z_j$. On the priority row,

$$\dot{q} = f_y(q), \quad f_y(q) = q - q^3 + \sum_{k \in \mathcal{R}_j} g_{jk}(y_k - q), \quad (\text{S2.5})$$

where $\mathcal{R}_j$ is the set of sources read into memory j and $g_{jk} = g(x_{jk})$ is the reader gain of link (j, k) . The same trajectory gives the trace inputs

$$h_{jk}^a = \left[S_a \frac{g_{jk}}{\kappa_r} (y_k - q)^2 - \epsilon_0 \right]_+, \quad h_{jk}^u = [S_u g_{jk} (y_k - q) f_y(q) - \epsilon_0]_+. \quad (\text{S2.6})$$

With $g = 0.7$, $y = 3$, a ten-unit reading, zero initial trace and a two-unit delay, the admission traces at pulse onset are 0.0439244 and 0.0506271 from $q(0) = +1$ and -1 , and the use traces are 0.0136507 and 0.1465084. The use ratio 0.0931735 therefore follows from the memory motion itself, with no fitted trace amplitudes.

Why alignment can still credit an uninformative source. Let $q^{(0)}(t)$ be the real trajectory without a weak non-valid reader of gain g_d , whose source value is an independent symmetric choice $y_d = \pm v_0$ with $|q^{(0)}| < v_0$. On the priority row, the leading use input averaged over the two signs is

$$\mathbb{E}_{y_d} \left[S_u g_d \dot{q}^{(0)} (y_d - q^{(0)}) - \epsilon_0 \right]_+ = \frac{1}{2} \left[S_u g_d |\dot{q}^{(0)}| \{v_0 - \text{sgn}(\dot{q}^{(0)}) q^{(0)}\} - \epsilon_0 \right]_+, \quad (\text{S2.7})$$

with a correction of order g_d^2 . An uninformative source value can therefore earn a positive trace by aligning with the memory motion. This mechanism is consistent with the 17 of 48 unrelated co-read links that use retained as false records when the hidden values were redrawn every cycle (Section S7.2). Those runs did not isolate the mechanism. Narrowing the reader lowers g_d and the unwanted input, but also changes acquisition. Positive alignment is not an unbiased measure of usefulness.

S3 Readout bounds and the functional window

This section supports the article's statements on sensing, protection and revision (article Fig. 4a–c and Methods Eqs. (16) and (17)). The results give sufficient conditions for specified record sets and initial preparations. They do not give the probability that learning reaches or remains in those sets.

S3.1 Sensing and reader gain at the two branches

During sensing the probe of the acting need is on, all $I_k = 0$, and every cavity keeps its holding drive. The sensing signal $D_k = \sum_j P_j |b_{jk}|^2$ is read at the end of the two-unit interval, and the source is chosen by the rule of article Methods Eq. (11) with $N_K = 3$, $\beta_s = 8$ and $\epsilon_x = 0.1$. At a fixed record and equal drive the high-to-low signal ratio is $D(x_L)/D(x_H) = 5.828427$.

For an ordinary reading the reader obeys article Methods Eqs. (8) and (9). After the fast transient, elimination of c_{jk} gives article Eq. (2),

$$\dot{z}_j|_{\text{readers}} \simeq \sum_k I_k K(x_{jk})(y_k - z_j), \quad K(x) = \frac{v^2}{\kappa_r + i\Delta_r(x)}, \quad g(x) = \text{Re } K(x) = \frac{g_0}{1 + [(x - x_H)/w]^2},$$

with $g_0 = v^2/\kappa_r$. The real part pulls the memory towards the source value; the imaginary part also acts on the complex memory. Assisted reading sets the selected link's detuning to zero, so its response is $K = g_0$ for that reading. Setting the source value to zero does not disconnect it. With $I_k = 1$ and $y_k = 0$, the reader still acts on the memory through $-Kz_j$. Disconnecting the source requires switching off both coupling terms.

At $w = 1$ the quiet gains are 0.70 and 0.0778. The weak tones alone change a high record by about 0.0067 in ten units at $\gamma = 0.001$. During repeated learning, active records can lie at 3.9 to 4.5, where the reader gain at $w = 1$ is about 0.57 to 0.32 and the measured sensing contrast is about 2 to 4. These quiet-branch values explain the mechanism; readout in the active network must be tested separately.

S3.2 Why learning needs assisted reading

All links start with low records. At $w = 1$ their gain is 0.0778, too small to reverse the memory against an opposing source value. The admission trace is also larger on unsuccessful unassisted readings than on successful ones (Section S5.2). The protocol therefore tunes only the selected reader of the acting need to resonance, with probability p_{as} , so that its source value enters at gain g_0 . At $p_{\text{as}} = 0.5$ all 16 valid links formed in the R1 learning test (T1), with median formation time 100.5 need-cycles. At $p_{\text{as}} = 0.2$ only one of 16 formed within 375 need-cycles. These are finite-horizon results at one working point, not a sharp threshold. Assisted reading also lets unrelated source values in: the observed median success at $p_{\text{as}} = 0.5$ was 0.87, against 0.9 for an action that follows a current valid source value.

S3.3 Comparison conditions for simultaneous readings

Let $V_{jk} \in \{0, 1\}$ be the assignment, with one valid source k_j per need; it is used only for evaluation. The functional target fixes a minimum valid-source selection probability p_* , a memory margin $m > 0$, a reading time T_{read} and the allowed set of simultaneous inputs. A formed record need not provide these functions. At $x = 7$, for example, $1/D(x) = 0.058824$ and, at $w = 1$, $g(x) = 0.050513$. Both are below the quiet low-record values 0.146447 and 0.077778.

Sensing with unequal competing records. Suppose the valid-source sensing intensity is at least $R_I > 1$ times every competing intensity. Scale the valid intensity to one and write each competing intensity as $u_l \leq 1/R_I$. Then $(1 - u_l)/(1 + \sum_{l'} u_{l'}) \geq (R_I - 1)/(R_I + N_K - 1)$ for every competitor l , so the selection rule gives

$$p_{jk_j} \geq p_{\min}(R_I, N_K) := \frac{\epsilon_x}{N_K} + \frac{1 - \epsilon_x}{1 + (N_K - 1) \exp[-\beta_s(R_I - 1)/(R_I + N_K - 1)]}, \quad (\text{S3.1})$$

which is article Methods Eq. (16). Equality holds when every competitor is at its upper bound, so the bound is tight for a given minimum contrast. Put $t_* = (p_* - \epsilon_x/N_K)/(1 - \epsilon_x)$ and $d_* = \beta_s^{-1} \ln[(N_K - 1)t_*/(1 - t_*)]$. For $1/N_K < t_* < 1$ and $0 < d_* < 1$ the sufficient contrast is

$$R_I \geq R_* := \frac{1 + (N_K - 1)d_*}{1 - d_*}; \quad (\text{S3.2})$$

if $d_* \geq 1$, no finite contrast meets the target. For record intervals $\mathcal{I}_+$ (valid) and $\mathcal{I}_-$ (non-valid), equal drive and settled fields give

$$R_{\text{box}} = \frac{\min_{x \in \mathcal{I}_-} D(x)}{\max_{x \in \mathcal{I}_+} D(x)}. \quad (\text{S3.3})$$

At $\delta = 3$, $N_K = 3$, $\epsilon_x = 0.1$, $\beta_s = 8$ and $p_* = 0.9$, $R_* = 3.9277463$. For $\mathcal{I}_- = [0.50, 0.75]$ the valid record must then satisfy

$$2.2627714 \leq x_{jk_j} \leq 3.7372286, \quad (\text{S3.4})$$

where the endpoints solve $D(x) = D(0.75)/R_*$; we call them the lower and upper sensing limits x_- and x_+ . They are operating limits, not the folds $x_{f\pm}$ of Section S1.1. For $\mathcal{I}_+ = [3.30, 3.60]$, $R_{\text{box}} = 4.4577206$ and $p_{jk_j} \geq 0.9091707$.

Protection depends on the total wrong gain. With $q = \text{Re } z_j$, use the comparison equation

$$\dot{q} = q - q^3 + \sum_{k \in \mathcal{R}_j} g_{jk}(y_k - q) + e_j(t), \quad |e_j(t)| \leq \eta_r, \quad (\text{S3.5})$$

where η_r bounds the residual of the real-memory equation. Assume $|y_k| \leq v_0$, $0 < m < q_{\max} < v_0$, an invariant outer bound $|q| \leq q_{\max}$, and total gain of the active non-valid readers at most $\bar{G}_d$ throughout the test. An answer that starts in $[m, q_{\max}]$ or $[-q_{\max}, -m]$ keeps its sign under non-valid-only readings if

$$\mathcal{M}_{\text{reject}} := m - m^3 - \bar{G}_d(v_0 + m) - \eta_r > 0, \quad (\text{S3.6})$$

which is article Methods Eq. (17). At $q = m$, the worst source values are all $-v_0$ and the worst residual is $-\eta_r$. The inequality makes the vector field point inward, ruling out a first outward crossing. The negative case follows by sign reversal. No independence of the non-valid inputs is needed.

Valid revision in finite time. Let the valid source value be $y_v = \pm v_0$ with reader gain at least $\underline{g}_v$, together with non-valid readers of total gain at most $\bar{G}_d$. In the coordinate $a = \text{sgn}(y_v)q$ the drift towards the valid answer is at least

$$F_{\text{rev}}(a) = a - a^3 + \underline{g}_v(v_0 - a) - \bar{G}_d(v_0 + a) - \eta_r, \quad (\text{S3.7})$$

because every non-valid source value is at least $-v_0$ in this coordinate and $|a| < v_0$. Every allowed initial answer reaches the valid sign with margin m within the reading if

$$F_{\text{rev}}(a) > 0 \quad (-q_{\max} \leq a \leq m), \quad T_{\text{rev}} := \int_{-q_{\max}}^m \frac{da}{F_{\text{rev}}(a)} \leq T_{\text{read}}. \quad (\text{S3.8})$$

Scalar comparison proves the time bound, and positive drift at $a = m$ prevents a return while the valid input continues. The minimum of the cubic is found from its endpoints and its interior stationary point $a = -\sqrt{(1 - \underline{g}_v - \bar{G}_d)/3}$ when $1 - \underline{g}_v - \bar{G}_d > 0$. The outer bound holds when

$$q_{\max} - q_{\max}^3 + G_{\text{tot}}^{\max}(v_0 - q_{\max}) + \eta_r < 0, \quad (\text{S3.9})$$

where $G_{\text{tot}}^{\max}$ bounds all active gains; for $q_{\max} = 1.8$, $v_0 = 3$, $\eta_r = 0.05$ and three gains of at most 0.7, the left side is -1.462 .

A failure boundary for simultaneous weak inputs. For constant total non-valid gain g_d , real memory, zero residual and all source values equal to $-v_0$,

$$\dot{q} = (1 - g_d)q - q^3 - g_d v_0. \quad (\text{S3.10})$$

For $0 < g_d < 1$ its positive maximum is $2(1 - g_d)^{3/2}/(3\sqrt{3}) - g_d v_0$, and the wrong-sign well disappears when this is negative. At $v_0 = 3$ the critical total gain is 0.108073762. One quiet low reader at $w = 1$ has gain 0.0777778, below the threshold; two have 0.1555556, above it. Starting from $q = 1$, the scalar model with two readers crosses zero at 3.89319. Integrating the full complex-reader equations gives 3.751196 (Section S7.6). Both crossings occur within the ten-unit reading while both records remain unformed. Here failure is proved directly, not inferred from a failed sufficient inequality. For one reader at $w = 1$, the eventual threshold lies at $x = 1.0739021$ and the ten-unit threshold near 1.2099938; both are below $x_U = 2$, so a single record can reverse an answer before it forms.

A nonempty common width range. Choose

$$\mathcal{I}_- = [0.50, 0.75], \quad \mathcal{I}_+ = [3.30, 3.60], \quad p_* = 0.9, \quad v_0 = 3, \quad m = 0.5, \quad q_{\max} = 1.8, \quad \eta_r = 0.05, \quad (\text{S3.11})$$

and allow up to two non-valid readers. Requiring $\underline{g}_v \geq 0.6$ and Eq. (S3.6) gives

$$\begin{aligned} w \geq w_{\min} &:= \frac{\max_{x \in \mathcal{I}_+} |x - x_H|}{\sqrt{g_0/0.6 - 1}} = 0.455081973 \dots, \\ w < w_{\max} &:= (x_H - 0.75) \sqrt{\frac{g_{\text{prot}}}{2g_0 - g_{\text{prot}}}} = 0.710092893 \dots, \quad g_{\text{prot}} := \frac{m - m^3 - \eta_r}{v_0 + m} = 0.0928571429 \dots, \end{aligned} \quad (\text{S3.12})$$

where g_{prot} bounds the total non-valid gain. Across this range, inserting $\underline{g}_v \geq 0.6$ and $\bar{G}_d < g_{\text{prot}}$ into Eq. (S3.8) verifies revision within ten units. The lower limit w_{min} follows from the convenience requirement $\underline{g}_v \geq 0.6$, which is sufficient for the revision-time bound but not necessary: a narrower reader with a smaller valid gain can satisfy Eq. (S3.8) directly, so the interval is a sufficient common range, not the set of all functional widths. The full-reader argument of Section S3.4 at $w = 0.6$ is a separate certificate and does not rest on this requirement. At $w = 0.6$, with $\bar{g}_d$ the largest gain of one non-valid reader with record in $\mathcal{I}_-$,

$$\begin{aligned} \underline{g}_v &= 0.6387564, & \bar{g}_d &= 0.0337891, & \bar{G}_d &= 2\bar{g}_d, \\ \mathcal{M}_{\text{reject}} &= 0.0884765, & T_{\text{rev}} &\leq 1.075396, & p_{jk_j} &\geq 0.9091707. \end{aligned} \quad (\text{S3.13})$$

The residual bound $\eta_r = 0.05$ is a premise here. The next subsection removes it by bounding the full reader, including its switch-on transient.

S3.4 The full-reader argument

The comparison above assumes a small residual from $t = 0$. We now bound that residual using the full complex reader and memory equations. Table S3 lists three applications of the argument and their results. All use $g_0 = 0.7$, $\kappa_r = 200$, $G = 80$ and $v^2 = g_0\kappa_r$. Readings are unassisted, with constant real source values satisfying $|y_k| \leq 3$ (valid value $y_v = \pm 3$). There is no correction drive, outcome writing or direct reader or memory noise. Time starts when the readers switch on.

Coordinates. Suppress the row index and write $z = q + ip$. On each active link put

$$r_k = -ivc_k, \quad L_k = \kappa_r + i\Delta_r(x_k), \quad K_k = \frac{v^2}{L_k} = g_k + ih_k, \quad (\text{S3.14})$$

so that during a reading

$$\dot{r}_k = -L_k r_k + v^2(y_k - z), \quad \dot{z} = (1 - |z|^2)q - i(G + |z|^2)p + \sum_k r_k. \quad (\text{S3.15})$$

Step 1: an invariant tube. For fixed records and zero initial readers, consider $|z| \leq R$ and $|r_k| \leq g_0(3 + R)$ for n active readers. On their boundaries

$$\frac{d|r_k|}{dt} \leq -\kappa_r|r_k| + v^2(3 + R) = 0, \quad \frac{d|z|}{dt} \leq R - R^3 + ng_0(3 + R), \quad (\text{S3.16})$$

so a negative second bound makes the product set invariant. Since $p(0) = 0$, the damping of the imaginary coordinate gives

$$|p| \leq p_R := \frac{ng_0(3 + R)}{G}, \quad (\text{S3.17})$$

and, for the radii used ($R^3 - R > 2/(3\sqrt{3})$),

$$|\dot{q}| \leq M_q := R^3 - R + ng_0(3 + R), \quad |\dot{p}| \leq M_p := (G + R^2)p_R + ng_0(3 + R), \quad |\dot{z}| \leq (M_q^2 + M_p^2)^{1/2}. \quad (\text{S3.18})$$

For states carried over from earlier dynamics and moving records, we propose a tube with $|q| \leq R$, $|p| \leq p_R$ and $|\dot{z}| \leq M$. Step 2 then proves that trajectories cannot leave it.

Step 2: the reference-response error. Define $e_k = r_k - \bar{K}_k(y_k - z)$, where $\bar{K}_k = \bar{g}_k + i\bar{h}_k$ is the response at a fixed reference population $\bar{x}_k$. With $L_k(t) = \kappa_r + i\Delta_r(x_k(t))$, the exact equation is

$$\dot{e}_k = -L_k(t)e_k + \bar{K}_k \dot{z} - i[\Delta_r(x_k(t)) - \Delta_r(\bar{x}_k)]\bar{K}_k(y_k - z). \quad (\text{S3.19})$$

No derivative of x_k appears, so the record path need only satisfy

$$|x_k(t) - \bar{x}_k| \leq d_x; \quad (\text{S3.20})$$

it need not be differentiable. For fixed records ($d_x = 0$, $\bar{K}_k = K_k$) the last term vanishes. Let S_K , H_K and G_K bound $\sum_k |\bar{K}_k|$, $\sum_k |\bar{h}_k|$ and $\sum_k \bar{g}_k$, respectively. Write S_d and H_d for the first two bounds restricted to non-valid readers, and put $R_z = (R^2 + b_p^2)^{1/2}$. Variation of constants gives, on the tube,

$$E(t) := \sum_k |e_k(t)| \leq E_0 e^{-\kappa_r t} + E_\infty, \quad E_\infty = S_K \left[\frac{M}{\kappa_r} + \frac{d_x}{w}(v_0 + R_z) \right], \quad (\text{S3.21})$$

where E_0 bounds the initial mismatch. For reader input carried over from earlier dynamics, take $E_0 = r_{\text{in}} + S_K(v_0 + \sqrt{R^2 + p_{\text{in}}^2})$, with $\sum_k v|c_k(0)| \leq r_{\text{in}}$ and $|p(0)| \leq p_{\text{in}}$. For zero readers and $z(0) = \pm 1$, take $E_0 = 4S_K$. The bound holds for each continuous noisy record path satisfying Eq. (S3.20). It assigns no probability to that condition. To close a proposed tube, combine it with the direct convolution bound $|\sum_k r_k(t)| \leq A_r(1 - u) + r_{\text{in}}u$, where $u = e^{-\kappa_r t}$ and $A_r = 3g_0(v_0 + R_z)$. For a reference envelope $B_r + E_0u$ with $0 \leq B_r < A_r$, the smaller of the two bounds never exceeds

$$\mathcal{E}(B_r) = \frac{(A_r - r_{\text{in}})B_r + A_r E_0}{A_r - r_{\text{in}} + E_0}, \quad (\text{S3.22})$$

the value at the intersection of the decreasing and the increasing affine function of $u \in [0, 1]$. Three checks establish the tube. At $q = \pm R$, the outward drift bound $R - R^3 + \mathcal{E}(G_K(v_0 - R) + H_K b_p + E_\infty)$ must be negative. The imaginary coordinate, with damping at least G and forcing at most $H_K(v_0 + R) + E(t)$, must remain below b_p . The resulting speed bounds must remain below M .

Step 3: the real residual. The real-memory equation then reads

$$\dot{q} = q - q^3 + \sum_k \bar{g}_k(y_k - q) + \xi(t), \quad |\xi(t)| \leq Rb_p^2 + H_K b_p + E_\infty + E_0 e^{-\kappa_r t}, \quad (\text{S3.23})$$

with b_p replaced by p_R in the fixed-record case. The whole reader-loading term is kept; no small residual is assumed at $t = 0$. The response constants S_K , H_K , $\bar{g}_v$ (the smallest valid reference gain) and $\bar{G}_d$ (the largest summed non-valid reference gain) are extrema of the Lorentzian response over the full record intervals. Their calculation includes resonance and the maximum of $|\text{Im } K|$ whenever either lies inside an interval. No trajectory sampling is used.

Step 4: protection during the reader switch-on. With non-valid readers only, let $E_{0,d}$ be the initial mismatch E_0 restricted to these readers and $\eta_{d,\infty}$ the bound of Eq. (S3.23) without its transient term. The transient part of ξ then has integral at most $E_{0,d}/\kappa_r$. Choose a startup band $[m, m + \delta_m]$ on which the non-transient worst-case drift $q - q^3 - \bar{G}_d(v_0 + q) - \eta_{d,\infty}$ is positive and increasing, with δ_m larger than this transient budget. A path from $|q(0)| \geq m + \delta_m$ to $|q| < m$ would have to cross the whole band after its last entry through $m + \delta_m$. Even with the positive drift discarded, the transient budget cannot carry it across, so the crossing is impossible. Sign reversal gives the negative case. The argument holds for every constant pattern of non-valid source values with $|y_k| \leq 3$.

Step 5: revision. After $t_0 = 0.05$ the transient is negligible. The tube already holds from switch-on, so at t_0 the memory is still in $[-R, R]$ and nothing is reset. In the coordinate $a = \text{sgn}(y_v)q$ the drift is at least $a - a^3 + \bar{g}_v(3 - a) - \bar{G}_d(3 + a) - \eta$, with η the post- t_0 residual bound. If this lower drift is positive on $[-R, 0.5]$, Eq. (S3.8) bounds the time needed to establish the valid margin 0.5. The bound is t_0 plus the integral. Positive drift at $a = 0.5$ then prevents a return.

Step 6: sensing on the same record set. With the probe on and readers off, each cavity has the reference field $\bar{b} = \tilde{f}/[1 - i(\delta - \bar{x})]$, and

$$|b(T_s) - \bar{b}| \leq |b(0) - \bar{b}|e^{-\kappa T_s} + d_x |\bar{b}|(1 - e^{-\kappa T_s}). \quad (\text{S3.24})$$

For $T_s = 2$, $\kappa = 100$ and $|b(0)| \leq 3$ the first term is below 7.0×10^{-87} . If ϵ_b bounds the right side divided by $|\bar{b}|$, the actual contrast is at least $R_{\text{box}}[(1 - \epsilon_b)/(1 + \epsilon_b)]^2$, and Eq. (S3.1) gives the sensing bound. Direct cavity-field noise adds a separate driven-field term.

Case 2 uses the fixed-record intervals

$$\mathcal{I}_+^{\text{fixed}} = [3.30, 3.70], \quad \mathcal{I}_-^{\text{fixed}} = [0.50, 0.75], \quad (\text{S3.25})$$

for which the largest valid denominator is $D(3.70) = 1.49$ and the smallest non-valid one $D(0.75) = 6.0625$.

Table S3: The three applications of the full-reader argument. Case 1 has one active reader; cases 2 and 3 allow one valid and up to two non-valid readers. Records are fixed in cases 1 and 2; in case 3 they may move by $d_x \leq 0.01$ (Eq. (S3.20)) around reference values in the stated intervals. All guarantees hold for ten-unit readings.

	Case 1: $w = 1$, one reader	Case 2: $w = 0.6$, fixed records	Case 3: $w = 0.6$, moving records
Records	low $[0, 0.75]$; high $[3.3, 3.6]$ (wide $[3.1, 4.6]$)	Eq. (S3.25)	non-valid $[0.50, 0.75]$; valid $[3.30, 3.60]$; $d_x \leq 0.01$
Initial state	$z(0) = 1$, $c(0) = 0$, $y = -3$	$z(0) = \pm 1$, $c_k(0) = 0$	$ q(0) \leq 1.8$, $ p(0) \leq 0.004$, $\sum v c_k(0) \leq 0.01$
Tube	$R = 1.8$; $ c \leq 0.283971830$; $ p \leq 0.042$; $ \dot{z} \leq 10.082038334$	protection $R = 2.1$, $p_R = 0.08925$, $ \dot{z} < 20.5$; revision $R = 2.5$, $p_R < 0.145$, $ \dot{z} < 34.5$	$R = 1.8$, $b_p = 0.05$, $M = 14$; closed with $ p < 0.048918$, $ \dot{z} < 13.516$
Key constants	low: $g \leq 0.086440735$, $ \xi \leq 0.025292468$ after 0.05; high: $g \geq 0.676644531$, drift bound 1.959160459	$\underline{g}_v \geq 0.570$, $\overline{G}_d \leq 0.0676$, $\overline{S}_K \leq 1.008$, $H_K \leq 0.572$, $S_d \leq 0.3076$, $H_d \leq 0.3001$	$S_K \leq 1.007587$, $H_K \leq 0.497859$, $G_K \leq 0.767579$, $\underline{\tilde{g}}_v \geq 0.638756$, $\overline{G}_d \leq 0.067579$
Protection	$\text{Re } z > 1/2$ throughout ($\dot{q} > 0.04716$ at $q = 1/2$)	$ \text{Re } z \geq 0.5$ on $[0, 10]$; residual $0.0751 + 1.2304e^{-200t}$, budget $0.006152 < 0.008$	sign kept with $ q \geq 0.5$ for $0.508 \leq q(0) \leq 1.8$; $\eta_{d,\infty} < 0.06565$, budget $< 0.007435 < 0.008$
Revision	$\text{Re } z = -1/2$ within 0.79043 (2.56695 on wide interval)	margin 0.5 from 2.745 (minimum drift 1.113228792, residual < 0.31)	margin 0.5 by 1.206 (1.156 for $d_x = 0$); residual < 0.181 (0.101)
Sensing	$p_H \geq 0.909170656$ (0.659201933 wide)	$p_v > 0.90289$ ($R_I = 4.068791946$)	$p > 0.90661$ (0.90917 for $d_x = 0$)

In case 1 the high-record bound uses $q(t) \leq 1 - (m_H - \eta_H)t + 0.014$, where g_H is the lower bound on the high-record gain, with $m_H = 3g_H - \frac{2}{3\sqrt{3}}(1 - g_H)^{3/2}$ and $\eta_H = 0.043742213$, so $T_{-1/2} \leq 1.514/(m_H - \eta_H)$. In case 2 the revision drift on $[-2.5, 0.5]$ is at least $F_{\text{fixed}}(a) = -a^3 + 0.3624a + 1.1972$, whose minimum $1.1972 - 2(0.3624)^{3/2}/(3\sqrt{3}) = 1.113228792 \dots$ gives $T \leq 0.05 + 3/1.1132 = 2.744933524 \dots$. In case 3 the tube closes with $R - R^3 + \mathcal{E}(\cdot) < -4.032 + 4.016210 < 0$, $E_0 < 4.847$ and $E_\infty < 0.151150$; the revision drift is $a - a^3 + 0.638756(3 - a) - 0.067579(3 + a) - \eta$, positive on $[-1.8, 0.5]$. Article Fig. 4b is covered by case 1 (one reader at $w = 1$) and by the failure boundary for two weak readers above; in article Fig. 4c, case 2 covers the protection trace from $z(0) = +1$, and case 3 with $d_x = 0$ covers the revision from $z(0) = -1.8$. Case 3 is a joint certificate for sensing and reading when the same reference populations serve both intervals. Case 2 establishes the readout capabilities of any record set in its interval. It does not widen the width range of Eq. (S3.12) or the allowed motion in case 3. None of these results shows that learning reaches or remains in the stated intervals.

S4 Retention and separation

This section supports the article's statements on retention of inactive records, revision of one need only, and equivalence of assignments. Priority acts through P_j , which selects the row that contributes to sensing and does not enter the reader gain $g(x_{jk})$. A change of priority can therefore redirect source selection without a change of the records. A change of validity is a different intervention: it changes which source reports a need's hidden value, while the links and the local rule stay the same, and later outcomes must erase obsolete records and form new ones (Section S7.4). A flip of a hidden value changes the correct answer, not the assignment.

S4.1 A finite residual-credit bound for dormancy

For admission and use, $P_j = 0$ sets the input of every trace on the dormant row to zero. A link that starts the hold with trace ϵ_d then has $\epsilon(t) = \epsilon_d e^{-t/T_\epsilon}$, even when other rows still receive outcome pulses. The total possible negative credit is therefore finite.

Consider the noiseless scalar record equation

$$\dot{x} = \gamma F_4(x) + R_+(t, x) + \Lambda \sigma(t) \epsilon_d e^{-t/T_\epsilon} H(x) Q(t), \quad R_+ \geq 0, \quad Q \geq -Q_{\max},$$

where R_+ allows extra nonnegative probe-tone and reader filling but no negative cavity-lag error, and $Q = s - \beta_{j^*} - \vartheta$ is the outcome signal of article equation (3). Let

$$\Psi(x) = x + x_* \ln(x/x_*), \quad \Psi'(x) = 1/H(x), \quad x > 0. \quad (\text{S4.1})$$

Choose a protected floor $x_{\text{fl}} \in [x_U, x_H]$ and start from $x_0 \geq x_{\text{fl}}$. The record never crosses below x_{fl} if

$$\Lambda Q_{\max} \int_0^\infty \sigma(t) \epsilon_d e^{-t/T_\epsilon} dt < \Psi(\min\{x_0, x_H\}) - \Psi(x_{\text{fl}}). \quad (\text{S4.2})$$

On $[x_{\text{fl}}, x_H]$, $F_4 \geq 0$, so in the Ψ coordinate all negative displacement is bounded by the left side. A first downward passage to x_{fl} would need at least the gap on the right, either from x_0 or after the last downward entry from above x_H . The strict inequality excludes it. If pulses last τ_w , start no sooner than a delay d_0 , and are separated by at least T_p , the total negative budget is at most

$$B_- = \Lambda Q_{\max} \epsilon_d T_\epsilon e^{-d_0/T_\epsilon} \frac{1 - e^{-\tau_w/T_\epsilon}}{1 - e^{-T_p/T_\epsilon}}. \quad (\text{S4.3})$$

For $\beta \in [0, 1]$ and $\vartheta = 0.2$ one may take $Q_{\max} = 1.2$. For a record initially at or above x_H , the reference parameters ($\Lambda = 10$, $T_\epsilon = 10$, $\tau_w = 2$, $T_p = 22$, $d_0 = 0$) give $B_- \simeq 24.5 \epsilon_d$, and the bounds $\epsilon_d < 0.0796$ and $\epsilon_d < 0.0501$ protect the formed branch and the floor $x_{\text{fl}} = 2.5$, respectively, in this noiseless scalar model; for a lower starting population the budget on the right of Eq. (S4.2) is smaller.

The bound explains why arbitrarily many later pulses cause only bounded erosion once the trace input is off. It does not turn the 400-outcome noisy tests into an indefinite stochastic guarantee: a full-field extension must bound any negative cavity-lag term in Ψ , and noise needs a separate first-loss calculation. Co-presence lacks the premise, because its trace on the dormant row still responds when another row reads a connected source.

S4.2 Assignment symmetry at fixed network size

For homogeneous, complete need–source wiring, apply a source permutation to the source values, all link states, reader controls and noise paths. Every link equation has the same parameters, the memory sum is unchanged by reindexing, and the selection probabilities permute with their inputs. The transformed trajectory therefore solves the same equations with the permuted assignment, and for identical initial laws and relabelled noise the observables have identical distributions. Any two injective assignments at fixed numbers of needs and sources are related in this way. This proves assignment equivalence for the homogeneous networks covered by D5. It does not guarantee learning in every history or bitwise equality under one random seed. The argument also does not cover heterogeneous links or other source counts. With all records equal and matched fast states, the same symmetry removes source preference and selective influence; whether an ordinary reading can still revise the answer is then set by the common population, and equal low records are too weak to do so (Section S3.2).

S5 Finite-pulse credit and the learning-retention target

The writer in article equation (3) updates records according to the outcome. The outcome sets the sign of a broadcast signal, and each link’s eligibility trace sets its share. This section explains why the update is centred on a baseline with a margin. It then derives the exact credit carried by a finite pulse and defines the joint target for learning and retention.

S5.1 The outcome baseline and the margin

A writer whose positive and negative updates are not balanced at the chance baseline can form records for useless sources: chance-level outcomes then carry net positive credit, and a record formed by chance is kept. The writer used in this paper is therefore centred on the chance baseline and carries a margin.

The centred writer with a margin. For a link (j, k) the outcome-dependent contribution to $\dot{x}_{jk}$ is

$$\Pi_{jk}(t) = \Lambda \sigma(t) \epsilon_{jk}(t) H(x_{jk}) Q, \quad Q = s - \beta_{j^*} - \vartheta, \quad H(x) = \frac{x}{x + x_*}. \quad (\text{S5.1})$$

Here $\epsilon_{jk} \geq 0$ is the local trace (Section S2) and $\Lambda > 0$ the writer gain. The pulse envelope is $\sigma(t) = 1$ for a duration τ_w after each non-guess action of the acting need j^* , and $\sigma = 0$ otherwise. The outcome s and the baseline β_{j^*} of the acting need are broadcast to every link. The baseline changes only on designated guesses (Section S1.3). The scale x_* is the population at which $H = 1/2$; it is not x_U . The same H acts on both signs, and its zero at $x = 0$ stops the writer alone from driving a record negative.

Hold x , ϵ and β fixed through a pulse and let the outcome have success probability ρ , independent of them within the event class. To first order the mean increment is $\Lambda \tau_w \epsilon H(x) (\rho - \beta - \vartheta)$. With $\beta = 0.5$ and $\vartheta = 0.2$ the feedback is $+0.3$ on success and -0.7 on failure, with mean $\rho - 0.7$: -0.2 at chance and $+0.2$ at $\rho = 0.9$. The margin therefore biases chance-level updates towards removal. Without it, a false high record whose outcomes sit at the baseline would be kept by the first-order mean drift.

The writer coordinate. For the writer alone, the coordinate $\Psi(x) = x + x_* \ln(x/x_*)$ of Eq. (S4.1) obeys $\dot{\Psi} = \Lambda \sigma \epsilon Q$, so a pulse with constant Q and area $\mathcal{A} = \Lambda \int \sigma \epsilon dt$ moves Ψ by exactly $\mathcal{A}Q$ and keeps $x > 0$. Centring is exact in Ψ : if $\mathbb{E}[\mathcal{A}Q] = 0$ the mean displacement in Ψ is zero, although the mean displacement in x is nonnegative because $(\Psi^{-1})'' = HH' > 0$. The baseline is not constant: with fair designated-guess outcomes and $\beta_0 = 1/2$, $\beta_{m+1} = (1 - \alpha)\beta_m + \alpha \hat{s}_m$ has mean $\frac{1}{2}$ and a limiting standard deviation of 0.0503 at $\alpha = 0.02$, and past baseline values shape the records and traces, so the covariance of the baseline with later traces cannot in general be set to zero. Section S5.2 gives the exact trace-weighted credit condition that replaces the first-order picture.

S5.2 Exact trace-weighted credit for finite pulses

Let Z denote the remaining state of the network (fields, readers, memories, traces, baselines and environment). Write one positive record as $\dot{x} = \mathcal{R}_x(t, Z) + \Lambda \sigma \epsilon H(x) Q$, where $\mathcal{R}_x = \gamma(|b|^2 - x + \chi|c|^2)$ in the network and $\gamma F_4(x)$ for the isolated quiet record. With $Y = \Psi(x)$ the chain rule gives

$$\dot{Y} = \frac{\mathcal{R}_x(t, Z)}{H(x)} + \Lambda \sigma \epsilon Q. \quad (\text{S5.2})$$

If $Q_n = s_n - \nu_n$ is constant during pulse n , with $\nu_n = \beta_{j^*} + \vartheta$, then over the complete event interval $[t_n, t_{n+1})$, which contains pulse n and ends at the start of the next event,

$$Y_{n+1} - Y_n = J_n + \mathcal{R}_n, \quad J_n = \mathcal{A}_n(s_n - \nu_n), \quad \mathcal{A}_n = \Lambda \int_{\text{pulse } n} \sigma \epsilon dt, \quad \mathcal{R}_n = \int_{t_n}^{t_{n+1}} \frac{\mathcal{R}_x(t, Z_t)}{H(x_t)} dt. \quad (\text{S5.3})$$

Equations (S5.2) and (S5.3) are deterministic chain-rule identities. For the noisy record of Section S1.5, the transformed dynamics has an additional drift while $x > 0$: the Itô correction $-D_\tau \gamma x_*/x^2$. It also has the noise term $\sqrt{2D_\tau \gamma} dW/H(x)$. Both contribute to $\mathcal{R}_n$, leaving the writer contribution J_n unchanged. These terms must therefore be included in the remainder when applying the identities to noisy histories. The contribution J_n remains exact even for large record changes during a pulse. Restoration during the pulse belongs to $\mathcal{R}_n$. Whenever the moments exist,

$$\mathbb{E}J = \mathbb{E}\mathcal{A}(\mathbb{E}s - \mathbb{E}\nu) + \text{Cov}(\mathcal{A}, s) - \text{Cov}(\mathcal{A}, \nu), \quad (\text{S5.4})$$

also conditionally on a pre-event state or an event class.

The exact mean-credit threshold. For fixed $0 < \nu < 1$ and $0 < \rho < 1$, with $\rho = \Pr(s = 1)$ and $\bar{\mathcal{A}}_s = \mathbb{E}[\mathcal{A} | s]$, if $(1 - \nu)\bar{\mathcal{A}}_1 + \nu\bar{\mathcal{A}}_0 > 0$ then

$$\mathbb{E}J > 0 \iff \rho > \rho_* := \frac{\nu\bar{\mathcal{A}}_0}{(1 - \nu)\bar{\mathcal{A}}_1 + \nu\bar{\mathcal{A}}_0}. \quad (\text{S5.5})$$

Proof. Conditioning on the score gives $\mathbb{E}J = \rho(1 - \nu)\bar{\mathcal{A}}_1 - (1 - \rho)\nu\bar{\mathcal{A}}_0$. Collecting terms in ρ gives the threshold. For $\mathbb{E}\mathcal{A} > 0$, this is article equation (18), $\mathbb{E}[\mathcal{A}s]/\mathbb{E}\mathcal{A} > \nu$. At $\nu = 0.7$, equal conditional areas give a threshold of 0.7. Failure-to-success area ratios of 3.8 and 4 raise it to 0.89864865 and 0.90322581. Here ρ is the scored action's success probability, not the receiving source's accuracy. An unrelated source read alongside a useful one shares its high-success broadcast. Its own joint distribution of area and outcome determines its credit. The threshold fixes the sign of mean writing. Formation and retention probabilities also depend on $\mathcal{R}_n$ and the order of events.

S5.3 The learning-retention target

A predictive criterion must connect acquisition to retention, and retention starts from the full learned state Z : records, fields, answer memories, traces, baselines, environment variables and the phase of the alternating schedule. The record pattern alone is insufficient. For assignment V , let $\mathcal{G}_V$ be the set of complete states whose records lie in the functional window of Section S3.3 and from which every allowed readout test succeeds. The learning-retention target is the probability of continuous function from the acquisition deadline t_A through an operating time t_R ,

$$\mathcal{C}_V = \Pr\{Z(t) \in \mathcal{G}_V \text{ for all } t \in [t_A, t_A + t_R]\},$$

with the score and feedback drawn jointly under the cycle law of Section S1.3 and the writer given no access to V . If a useful state is reached early and then lost, it counts only if it is recovered by the deadline and retained thereafter. For fixed records in the certified range and the initial preparations specified in Section S3.4, the corresponding readout tests are guaranteed; the record window alone does not fix the initial reader transient. The sampled monitors used in the network studies estimate an operational version of $\mathcal{C}_V$, from sensing at cycle ends and the readings that occurred; they do not certify continuous retention, which would require bounds on the paths between samples. The G3 durability target permits loss and recovery, unlike strict retention.

S6 The self-limiting writer

The original writer can carry a formed valid record beyond the sensing limit x_+ . Positive writing continues after formation, while the restoring flow above x_H is weak (Sections S7.7 and S7.8). This section derives conditions for a local writer change to preserve the readout it creates. We change the positive branch while retaining the negative branch, traces and controls. Section S7.10 tests the resulting law.

The law and its physical reading. Article equation (4) replaces article equation (3) by

$$\Pi_{jk} = \sigma \epsilon_{jk} H(x_{jk}) [\Lambda_+ f_w(x_{jk}) [Q]_+ - \Lambda_- [-Q]_+], \quad f_w(x) = a_w(x) \kappa_{\text{cap}}(x), \quad Q = s - \beta_{j^*} - \vartheta, \quad (\text{S6.1})$$

$$a_w(x) = \frac{1}{1 + [(\delta_w - x)/w_w]^2}, \quad \kappa_{\text{cap}}(x) = \begin{cases} 1, & x \leq x_a, \\ (x_c - x)/(x_c - x_a), & x_a < x < x_c, \\ 0, & x \geq x_c. \end{cases} \quad (\text{S6.2})$$

The original writer has $f_w \equiv 1$ and $\Lambda_+ = \Lambda_- = \Lambda$, corresponding to non-resonant injection without saturation. The factor $a_w(x)$ describes absorption of a write tone at normalised detuning δ_w through the resonance shifted by the record. Its width is $w_w = 1$ in the units of Eq. (S1.4). The factor $\kappa_{\text{cap}}(x)$ describes saturation of a positive channel whose available states are exhausted at x_c , with a linear taper from x_a . The design choices are δ_w , peak positive gain Λ_+ and interval $[x_a, x_c]$. The negative branch retains gain Λ_- at every record value, allowing erasure for replacement. Both factors act on every link and depend only on its population, without information about validity. The rate is continuous and locally Lipschitz. It neither clips nor resets x . Above x_c , positive writing stops while restoration and negative writing continue.

Physical origin of the two factors. The absorption line follows from the element's field response. Let a write tone of normalised power A_w and detuning δ_w drive the same cavity mode as the holding drive at $\delta = 3$. At fixed x the field equation is linear, so the settled field is the sum $b = b_0 + b_w e^{i\Omega t}$ of the single-tone responses, with $|b_0|^2 = A_0/D(x)$, $|b_w|^2 = A_w/[1 + (\delta_w - x)^2]$ and beat frequency $\Omega = \kappa(\delta - \delta_w)$. If the write envelope and the reservoir vary slowly compared with Ω , the reservoir responds to the averaged intensity

$$\langle |b|^2 \rangle \simeq \frac{A_0}{D(x)} + \frac{A_w}{1 + (\delta_w - x)^2},$$

and the interference term drives only a population oscillation of amplitude about $2\gamma|b_0||b_w|/\Omega$, of order 10^{-4} in x at the reference parameters. Matching the filling rate $\gamma A_w a_w(x)$ to the line term requires $\gamma A_w = \Lambda_+ \sigma \epsilon H(x) [Q]_+$; with the low-record trace areas measured below, the equivalent write power during a pulse is of order 60 near the barrier top (about 40 at x_L), compared with $A_0 = 4$. The cavity relaxation time $1/\kappa = 0.01$ is much shorter than the two-unit pulse and the reservoir changes little during a beat, so the averaged description holds away from the pulse edges; a device must also avoid saturating the cavity or the shared reservoir with the write tone. The derivation does not supply the factor $H(x)$ or the per-link control of write power by trace and outcome, which are retained from the original writer as proposed device operations. The capacity is phenomenological: saturating the shared reservoir would change the holding inflow and the

bistability, so κ_{cap} requires a separate filling channel whose occupation follows x adiabatically, and a slow channel would be an additional state variable. The linear taper between x_a and x_c is a modelling choice.

The capacity boundary. The negative branch is additive in Ψ . For the positive branch define $\Phi(x) = \int^x du/[H(u)f_w(u)]$; for the writer alone $\dot{\Phi} = \Lambda_+ \sigma \epsilon[Q]_+$, so a positive pulse moves Φ by exactly $\Lambda_+ (\mathcal{A}/\Lambda)[Q]_+$ however much the record changes. With κ_{cap} present, $\Phi \rightarrow +\infty$ logarithmically as $x \rightarrow x_c^-$: the record approaches x_c exponentially in the dose, and no finite positive writer-only pulse can reach or cross the boundary from below. This result is exact for finite pulses. Confinement of the full dynamics is a separate question. With $\dot{x} = \mathcal{R}(t, Z) + H(x)[\Lambda_+ f_w(x)u_+ - \Lambda_- u_-]$, $u_{\pm} = \sigma \epsilon[\pm Q]_+$ and $\mathcal{R} = \gamma(|b|^2 - x + \chi|c|^2)$, confinement to $x \leq x_c$ also requires the passive flow to point inward at x_c , that is $|b|^2 + \chi|c|^2 \leq x_c$ there. In the fast-field reduction at the largest weak-tone power the available margin is

$$x_c - \frac{4.8}{D(3.60)} = 0.0705882\dots, \quad (\text{S6.3})$$

which must accommodate reader loading and reduction error; the full fields do not automatically satisfy it (a state with $|b|^2 = 4.8$ at $x = 3.60$ has outward drift 0.0012), although during the writer pulse the weak tones are off and the inflow $4/D(x_c) = 2.94$ is well below x_c . Noise removes the strict boundary, because additive reservoir noise is unbounded: operating valid records under the line with the capacity reach 3.64 in the development study (Section S7.10) and 3.63–3.66 in the six-network validation (Section S7.11). The line alone has $f_w(x_c) = a_w(x_c) > 0$ and confines only through balance between positive input, negative writing and restoration; with the capacity, positive writing vanishes at x_c , so the passive-flow condition at the boundary no longer imposes an upper bound on the positive writer gain. Below x_a the capacity leaves the positive branch unchanged, and the negative branch is unchanged for every f_w : large negative pulses are allowed deliberately, since relearning by erasure requires some negative sequences to remove a record.

Mean-flow design screen with measured event classes. Estimate the mean writer drift of Ψ per need-cycle by a sum over event classes with rate r per need-cycle, mean trace area per unit writer gain $\bar{a}$, success probability ρ and mean outcome signals $q_+ = \mathbb{E}[Q | s = 1]$ and $q_- = \mathbb{E}[-Q | s = 0]$:

$$\bar{J}(x) \simeq \sum_{\text{classes}} r \bar{a} [\rho q_+ \Lambda_+ f_w(x) - (1 - \rho) q_- \Lambda_-]. \quad (\text{S6.4})$$

The estimate makes two approximations: within each class it treats the trace area and the outcome as independent, using products of class means in place of the class mean of $\mathcal{A}Q$ (the conditional-area threshold of Section S5.2 shows that this covariance shifts the acquisition threshold), and it evaluates the population-dependent factors at the current population rather than along the pulse. The quiet restoring drift of Ψ per need-cycle is $R(x) = \gamma T_{\text{ev}} F_4(x)/H(x)$, with $T_{\text{ev}} = 44$. The screen asks three things. For *acquisition*, the mean flow must cross the formation barrier,

$$\bar{J}_{\text{low}}(x) + R(x) > 0 \quad \text{for all } x \in [x_L, x_U]. \quad (\text{S6.5})$$

For *confinement*, the high-record mean flow must have a stable zero x_{eq} below the upper sensing limit x_+ for the specified competitor range, with a margin that exceeds record fluctuations and any downward shift of x_+ by transient competitor records; a capacity with $x_c < x_+$ supplies a deterministic margin under the boundary conditions above. For *replacement*, an obsolete record read at the high-record rate and scored at chance must satisfy $\max_{x \in [x_U, x_c]} [\bar{J}(x) + R(x)] < 0$; with an unshaped positive branch the maximum lies at the erasure barrier of the restoring flow, $x = 2.857$, and when f_w varies the maximum moves. The event classes were measured in the original-writer arm of the development study at $w = 0.6$ and $\Lambda = 5$, using every cycle of the 1250-cycle histories:

Record band	Class	r	$\bar{a}$	ρ	q_+	q_-
$x_v < 2$	assisted valid reading	0.201	0.0827	0.901	0.296	0.709
$x_v < 2$	unassisted valid reading with trace	0.024	0.0276	0.115	0.296	0.709
$3 < x_v < 3.8$	valid reading	0.719	0.0781 (success), 0.0782 (failure)		0.911	0.296

These are admission statistics. The corresponding use statistics (assisted valid readings at $x_v < 2$: $r = 0.240$, $\bar{a} = 0.0855$, $\rho = 0.910$; valid readings at $3 < x_v < 3.8$: $r = 0.732$, $\bar{a} = 0.0744$ after success and 0.0786 after failure, $\rho = 0.899$) give the use values quoted below. The screen is sensitive to such details: separating success and failure trace areas in the low band lowers the use acquisition threshold from 5.03 to 4.8.

The original writer cannot satisfy both conditions at $w = 0.6$. With $f_w \equiv 1$ and the admission statistics, Eq. (S6.5) requires $\Lambda \geq 7.36$, whereas keeping $x_{\text{eq}} \leq x_+ = 3.73723$ with non-valid records at 0.75 requires $\Lambda \leq 5.50$, and confinement to $x_{\text{eq}} \leq 3.70$, the upper edge of the proved reader window (Section S3.4), requires $\Lambda \leq 4.9$. No Λ satisfies both acquisition and confinement in this screen. At the tested value $\Lambda = 5$ the acquisition margin is -0.0068 per need-cycle, making formation slow and fluctuation-driven, and the equilibrium $x_{\text{eq}} = 3.706$ lies only 0.031 below x_+ , so finite pulses, noise and moving competitor records can overcome the margin (Section S7.7). With the use statistics the admissible range is $5.03 \leq \Lambda \leq 6.17$, narrowing to $5.03 \leq \Lambda \lesssim 5.5$ for $x_{\text{eq}} \leq 3.70$, which places $\Lambda = 5$ near its lower edge and is consistent with use giving faster but equally short-lived acquisition. A uniformly stronger writer meets acquisition but fails confinement ($\Lambda = 8.75$: margin $+0.0040$, $x_{\text{eq}} = 3.95$); a uniformly weaker one does the reverse ($\Lambda = 1.35$: margin -0.0173 , $x_{\text{eq}} = 3.49$).

How the two factors divide the work. The table compares the screen's three quantities for the laws tested in the development study, with $\Lambda_- = 5$ and the admission statistics (use values in parentheses): the acquisition margin is the minimum over $[x_L, x_U]$ of the valid record's net flow $\bar{J} + R$, x_{eq} its stable high equilibrium, and the last column the maximum over $[x_U, x_c]$ of the obsolete record's net flow.

Realisation	Λ_+	$f_w(x)$	Acquisition margin	x_{eq}	Max. obsolete flow
original	5	1	-0.0068 (-0.0001)	3.706 (3.674)	$+0.005$
line	8.75	$a_w, \delta_w = 1.25$	$+0.0090$ ($+0.0186$)	3.441 (3.423)	-0.016
capacity, original gain	5	$\kappa_{\text{cap}}, [3.30, 3.60]$	-0.0068 (-0.0001)	3.479 (3.470)	$+0.005$
capacity, peak gain	8.75	κ_{cap}	$+0.0099$ ($+0.0205$)	3.515 (3.508)	$+0.036$
line with capacity	8.75	$a_w \kappa_{\text{cap}}$	$+0.0090$ ($+0.0186$)	3.403 (3.390)	-0.016
same gains, no shape	8.75	1	$+0.0099$ ($+0.0205$)	4.030 (3.975)	$+0.036$
line on both branches	8.75 (both)	a_w (both)	$+0.0038$ ($+0.0149$)	3.498 (3.490)	$+0.036$

The absorption line governs acquisition and replacement: it is strong where the formation barrier is crossed ($a_w(1.38) = 0.983$) and weak between x_U and x_H ($a_w(2.86) = 0.28$), so the predicted mean flow erases an obsolete record without relying on diffusion. The capacity sets the upper boundary: with $x_c = 3.60 < x_+$, the line's finite margin ($\Psi(x_+) - \Psi(x_{\text{eq}}) = 0.379$, about 10.9 mean success kicks) becomes a boundary for positive writing, and deterministic flow returns a record carried above x_c by noise. The last two rows separate the law's shape from its gains: the unshaped writer with the line's asymmetric gains acquires, but its fixed point lies above x_+ and its flow at the erasure barrier is positive, so the screen predicts excessive detuning and failed replacement; the line on both branches acquires and confines but is predicted to retain an obsolete record. At fixed peak gain, a write line centred far below the barrier top ($\delta_w \lesssim 0.9$) or above it ($\delta_w \gtrsim 1.7$) loses acquisition because a_w is then too small at the barrier. Figure S3 shows the region of the $(\delta_w, \Lambda_+/\Lambda_-)$ plane satisfying acquisition and $x_{\text{eq}} \leq 3.55$.




Figure S3: **Mean-flow design map.** Admission rule at $w = 0.6$ and $\Lambda_- = 5$, in the plane of write-line detuning δ_w and peak positive gain Λ_+/Λ_- . Solid contour: the acquisition condition, Eq. (S6.5). Dashed, dash-dotted and dotted contours: the stable fixed point x_{eq} of the high-record mean flow at 3.5, 3.6 and the upper sensing limit $x_+ = 3.737$. The shaded region satisfies acquisition and $x_{\text{eq}} \leq 3.55$, a design margin below x_+ . P marks the tested line (Section S7.10). The map omits the capacity boundary $x_c = 3.60$ and uses event statistics measured under the original writer; it is a design screen, not a bound.

Why the negative branch is left unchanged. Replacement requires negative mean credit to exceed the upward restoring flow between x_U and x_H . Applying the same line to both branches multiplies the negative credit at the erasure barrier by $a_w(2.86) = 0.28$: with $\Lambda_{\pm} = 8.75 a_w(x)$, net writer credit at the barrier is -0.028 against a restoring flow of $+0.063$, so an obsolete record is retained near $x = 3.3$ and the screen predicts failed replacement. The development study confirms this under admission (Section S7.10); under use, the large failure pulses do erase the record, but seven times more slowly than the line. A second requirement is that low non-valid records have no net upward mean flow. An unassisted reading whose non-valid source opposes the stored answer is scored correct with the stale-answer probability $0.7048 = (1 + 0.8^4)/2$, within the baseline fluctuation of $\nu = \beta + \vartheta = 0.7$, so the class has approximately zero first-order mean credit under the original writer and small positive credit whenever $\Lambda_+ f_w(x) > \Lambda_-$ at the non-valid record (the tested line gives 1.21 at x_L and 1.12–1.40 over $[0.50, 0.75]$). Its trace areas are small ($\bar{a} \simeq 0.02$) and its drift is an order of magnitude below the low-branch restoring flow, leaving the total flow inward; for assisted non-valid readings ($\rho = 0.556$), assuming equal mean trace areas after success and failure, mean credit is negative if $\Lambda_+ f_w(x)/\Lambda_- < (1 - \rho)q_-/(\rho q_+) = 1.91$, which the tested ratios satisfy. This class limits how far the peak positive gain can rise: the one persistent false record among the development arms appeared at $\Lambda_+ = 10$.

Scope. The boundary and coordinate results are exact for the deterministic equations under the stated premises. The acquisition, confinement and replacement conditions are mean-flow screens, using products of class means rather than the class mean of $\mathcal{A}Q$, with statistics measured under the original writer; the absence of a Λ satisfying both conditions is a conclusion of the screen, which the original, weak and strong control arms test directly (Section S7.10). Margins expressed in mean kicks guide the design; they are not bounds on exit probability.

S7 Simulation evidence

This section follows the article’s claims: study design and registered criteria (Section S7.1), learning and false credit (Section S7.2), transplant (Section S7.3), retention, revision and continued use (Section S7.4), robustness (Section S7.5), and the functional window on stored states (Section S7.6). The studies at $w = 0.6$ follow.

S7.1 Design, statistics and criteria

Trace-rule studies. The trace-rule evidence comprises the historical tests R1–R5, the fresh G2 cohort and its G3 continuation. G2 compares the three trace rules on 24 new, matched networks; G3 continues those histories to 3000 decision cycles and adds interventions. The historical test sets R1–R5 (eight seeded networks per condition, criteria fixed before the runs, 62 registered statements) are reported here only where the article uses them: the eight-pair transplant of article Fig. 2c,d (R1), the permuted-outcome control (R4) and the robustness tests of Section S7.5 (R3, R5). The G1 audit replayed and rescored them (Section S9.2). The complete registry, including the statements that failed, will be deposited (Section S9). We report the sets separately, without pooling them, and report each trace rule separately.

Units and statistics. The sampling unit is one seeded network. Links, checkpoints and broadcast outcomes within a run are clustered observations. Network counts have exact binomial intervals, and link rates have intervals obtained by resampling whole networks. Comparisons between rules use paired-network resampling. In G2 and G3, count comparisons use an exact paired test when the arms share random streams and Fisher’s exact test otherwise; the writer studies state their own pairing (Sections S7.11 and S7.13). In G2 and G3 each random stream (environment, source choice, assisted reading, guesses, noise, assignment and diagnostics) is separate and hash-labelled. Different networks have disjoint streams. Variants and intervention arms of the same network share their external streams, which pairs the comparisons without forcing adaptive trajectories to agree.

Measures. A record is formed when $x_{jk} > x_U = 2$. It is persistent when its logged value stays above this threshold for more than 20 consecutive decision cycles. Formation at the end of a phase is scored separately. Formation time, called acquisition time in the registered criteria, is the first need-cycle at which a valid record crosses upward. Erasure is the first downward crossing after an assignment change. A clean network ends with both valid records formed and no other record formed. R1 and R2 late-phase measures use the last 75 decision cycles of a phase; R5 uses the last 75 need-cycles.

S7.2 Learning and false credit

G2 uses the reference parameters and one executable in which only the trace rule changes ($S_u = 1$ for use). Its 24 networks cover all six assignments of two valid sources among three, each with four initial hidden-value pairs (12 same-sign and 12 opposite-sign). No historical seed is reused. Learning runs for 600 + 150 decision cycles and the null for 600. The step is $\Delta t = 0.002$, with matched refinements to 0.001 and, where required, 0.0005 (Section S9.1). Formation passes if both valid records are formed at the end in at least 21/24 networks with at most 9/96 persistent non-valid links. The co-reading test passes with at most 3/48 persistent unrelated co-read links and is conclusive only if at least 18/24 networks keep both valid records. The null passes with at most 12/144 persistent links.

Table S4: G2 results at the reference step. A star marks a count with unresolved numerical components at the trajectory level; the criterion verdicts do not change over the tested steps. The markers [†] and [‡] are explained in Section S9.1.

Test or measure	C: co-presence	A: admission	U: use [†]
Both valid formed, separate reading	13/24	23/24	23/24
Persistent non-valid links, separate reading	34/96*	0/96	0/96*
Formation criterion	FAIL	PASS	PASS
Persistent null links	12/144* (at the threshold)	0/144	2/144
Null criterion	PASS	PASS	PASS
Persistent co-read links, flip protocol	33/48*	41/48	2/48
Both valid formed, flip co-read	12/24*	8/24*	22/24
Flip co-read criterion	INCONCLUSIVE	INCONCLUSIVE	PASS
Persistent co-read links, redraw protocol	35/48*	47/48	17/48
Both valid formed, redraw co-read	24/24*	12/24	24/24
Redraw co-read criterion	FAIL	INCONCLUSIVE	FAIL

Admission and use form the valid records selectively; co-presence fails the formation criterion. Exact 95% Clopper–Pearson intervals are [0.79, 1.00] for 23/24 and [0.33, 0.74] for 13/24. The flip co-read arm is

inconclusive for admission and co-presence because fewer than 18 of 24 networks kept both valid records (8 and 12). The counts 41/48 and 33/48 describe those arms but carry no registered verdict. Use reduces co-read false credit under the flip protocol, but the redraw result shows that this reduction depends on the task. Median formation times are 86.5 need-cycles for admission and 67.5 for use. Networks that had not formed by the end were counted at the cap, which leaves the median unchanged.

With the writer’s outcomes permuted, the persistent non-valid counts are 77/96 for C, 38/96 for A and 68/96 for $U^\dagger$. The permutation follows the procedure of the historical R4 control: for each need, a seeded permutation of that network’s own scored non-guess outcomes from its native run, drawn from an independent stream, with the true scoring, baseline and correction retained. Learning therefore depends on the timing of trace and outcome, not only on a positive mean score. The G2 null uses an independent unrelated variable for each source and is not pooled with the shared-variable historical null. In co-read tests, the additional source’s record often formed after the valid record; its later readings then corrupted the answer, causing failed actions followed by erasure and re-formation of both records. Under use, the pulse-time trace of the co-read source was about a tenth of the valid source’s (0.086 flip, 0.099 redraw).

S7.3 Transplant

Two networks are trained on exchanged assignments for 1000 decision cycles each. Their answer memories are set to $(+1, +1)$, their traces to zero, and their fast fields are allowed to relax for 50 time units. Only the six records are exchanged. The full state then relaxes for another 50 quiet units. Each source is read with both signs of its value for ten units with the probe off, followed by two units with the readers disconnected. Influence is the largest change of $\text{Re } z_j$ over the two signs. Sensing is measured separately under each need’s probe. The equal-record control sets all records to their mean.

Historical pairs (R1; article Fig. 2c,d). The criteria were fixed before the runs: the sensing argmax and the influence must follow the exchanged records in at least seven of eight pairs, and the equal-record control must remove both in at least seven of eight. The sensing preference followed the exchanged pattern in all 16 network–need pairs of the eight admission pairs. The selected signal exceeded the strongest competing signal by a factor 2.35–4.0. The pattern’s valid source changed its memory by 2.00, while changes in the other memories were at most 0.019. With equal records, every source could change both memories by about 2.00: selectivity was lost, but memory revision remained possible. Before the exchange, both valid records had formed in all eight recipient networks.

Fresh pairs (G2). The transplant follows the exchanged records in 4/24 C pairs, 21/24 A pairs and 15/24 $U^\dagger$ pairs at the reference step. Equalising the records removes selectivity in 22/24 (C), 24/24 (A) and 22/24 (U) pairs. Exact 95% intervals are $[0.68, 0.97]$ for 21/24, $[0.41, 0.81]$ for 15/24 and $[0.05, 0.37]$ for 4/24. For U, both networks had formed records before the exchange in 16 pairs, and 15 of those transferred. At the refined step, the A count is 20 and the C count 5. No numerical pass threshold was fixed for these fresh assays.

S7.4 Retention, revision and continued use

G3 continues the G2 histories to 3000 decision cycles (1500 per need, or 66 reservoir lifetimes). The first 750 rows are identical to G2, so G3 is not an independent cohort. Parameters are unchanged, and all 24 networks remain in every denominator.

Dormancy. In the fresh holds, each G2 network’s 750 learning cycles are followed by 400 cycles in which only need I acts, then 20 cycles of need E. A hold passes if the held record stays above 2.5 throughout and the valid-source fraction on reactivation is at least 0.6. Fresh dormancy succeeds in 0/24 C, 23/24 A and 24/24 $U^\dagger$ networks; all 23 A records formed at entry pass, and the remaining A network entered without a formed record. The residual-credit bound of Section S4.1 explains this retention. Holds under Poisson record noise met the criterion in five of eight networks at $\Omega = 12$ (Section S7.5).

Durability. Let O_f count the rows in cycles 2501–3000 with both valid records above x_U and every non-valid record at or below it. A network is durable if $O_f \geq 400$ and no non-valid link has a persistent false record in cycles 1–3000; the target is at least 21 of 24. This measures late correct-only occupancy and cumulative false credit, and it allows earlier loss and recovery.

Table S5: G3 continued operation at the reference step. Stars mark unresolved component counts; durability verdicts agree at all three tested steps.

Measure	C [‡]	A	U [†]
Durable networks	0/24; FAIL	24/24; PASS	21/24; PASS, at the target
Persistent non-valid links, 1–3000	73/96*	0/96	1/96*
Both valid formed at cycle 3000	16/24*	24/24	23/24*
Mean correct-only fraction, cycles 2501–3000	0.440*	1.000	0.978

Admission stays in the correct-only pattern throughout cycles 2501–3000 in every network ($O_f = 500$ in all 24). Use loses and recovers records but reaches high late occupancy. Only three networks are discordant in the A–U durability comparison (exact paired test, $P = 0.25$), so it does not establish a difference in durability probability. The pre-run prediction that use would fail durability was wrong: an early loss estimate and a slower assumed reacquisition understated its later occupancy (article Fig. 3a,b).

Reader withdrawal. At checkpoints 2000 and 3000, paired branches start from the same learned state and run for 200 cycles with assisted reading off. One keeps the readers; the other disconnects them. Both retain the writer and correction drive. Every action is scored at $t = 14$, before correction. Success is counted over all actions, including designated guesses. If every non-guess action copied a current valid reading, the flip protocol would give $0.2 \times 0.5 + 0.8 \times 0.9 = 0.82$; this, not perfect prediction of a changing hidden value, is the reference for the live-reader values near 0.80.

Table S6: Paired reader withdrawal. Δ is live minus blocked success over all actions; LB is its one-sided 95% paired-network lower bound; the last column excludes designated guesses. All six differences are positive and stay positive under the tested numerical alternatives.

Variant	Checkpoint	Δ , all actions	LB	Δ , decided actions
C [‡]	2000	0.0365	0.0273	0.0454
C [‡]	3000	0.0448	0.0354	0.0555
A	2000	0.1383	0.1302	0.1715
A	3000	0.1383	0.1304	0.1720
U [†]	2000	0.1342	0.1248	0.1661
U [†]	3000	0.1346	0.1256	0.1673

With readers blocked, success is about 0.666–0.668 over all actions and 0.704–0.708 over decided actions. Designated guesses contribute chance success to the first value. An eight-network diagnostic also removes the correction drive. With live readers, the correction’s contribution to decided-action success is 0.0047 for A (interval includes zero) and 0.025 for U[†]. With both readers and correction removed, success is near 0.52.

One-row change. At cycle 2000, need E’s valid source is exchanged with the unrelated source while need I is unchanged. The original assignment returns at cycle 2751. Each changed branch is paired with an unchanged continuation. For A, the obsolete record is formed at reassignment in all 24 networks; the Kaplan–Meier median erasure time is 21 E-need-cycles and the median new-formation time 156. For U[†], 23 networks enter with the obsolete record formed, and the medians are 6 and 87. Neither variant shows excess sustained losses on the unchanged row relative to its paired continuation. Return-time comparisons have intervals that span zero. The one-row change is a small version of the sequential-learning test for catastrophic interference [6, 7]: one row must be relearned without disturbing the other. Record revision is slower than the reader’s response. The reader receives no direct signal that the assignment has changed; adaptation requires later outcomes.

S7.5 Robustness

These are historical tests of the original writer at $w = 1$ with their own registered criteria. They show that record formation and the trace-rule results do not depend on one parameter point; they do not validate robustness of the final six-criterion learning sequence, which was tested only at the parameters of Sections S7.11 and S7.13. R3 used 184 runs of the original code: a one-parameter-at-a-time scan at 14 settings with four networks each (not a full parameter box; at changed δ the quiet fixed points were solved numerically and the reader centred at the new x_H) and count-noise tests with Poisson increments at occupation scale Ω

(Section S1.5); its runs of a second implementation are described in Section S9.2. R5 used 152 runs with the use trace (article equation (14), $S_u = 1$), including slower fields and a three-need, four-source network with a randomly assigned valid source for each need and a fourth source valid for no need; its baseline variant is in the registry.

Table S7: Robustness tests (R3 and R5), with the registered criteria in brief; the full criteria are in the registry to be deposited.

Condition	Result
14 one-parameter settings (R3)	formation and rejection at 14 of 14 settings with no false record; replacement at 14 of 14
Poisson increments, $\Omega = 50$	formation, null and dormancy criteria hold
Poisson increments, $\Omega = 12$	rejection degrades: 6 of 32 persistent non-valid records; null 1 of 48; dormancy 5 of 8
Three needs, four sources (R5)	all valid records formed in 6 of 8; 3 persistent false records of 72; replacement under a cyclic shift in 5 of 8; null 4 of 96; co-read 2 of 24 but only 4 of 8 runs keep all valid records (inconclusive)
Slower fields (R5)	220 and 55 cavity amplitude-relaxation times per cycle (reference 2200): both end with all valid records in 6 of 8 and no persistent false record of 32

The slowest field setting has $\kappa/\gamma = 2500$; this is the smallest separation tested, not a lower bound for every implementation. The three-need formation count of six of eight sits at the declared threshold.

S7.6 The functional window on stored states

Frozen-record reader tests (article Fig. 4b,c). Each test fixes the reservoir populations, starts all reader fields at zero and holds the listed real source values for ten time units. It integrates the full complex reader and memory equations with $G = 80$, $\kappa_r = 200$, $g_0 = 0.7$, $v^2 = g_0\kappa_r$ and $u_j = 0$; there is no record evolution, outcome writer, assisted reading or noise. Two quiet low-record readers reverse the answer although

Table S8: Frozen-record reader tests. Each comma-separated entry is one active reader. The first zero is the first crossing of $\text{Re } z$; “none” means no crossing in ten units.

Records	Source values	w	$z(0)$	$\text{Re } z(10)$	First zero
x_L	-3	1.0	+1	0.789827	none
x_L, x_L	-3, -3	1.0	+1	-1.124045	3.751196
x_L, x_L	-3, -3	0.7	+1	0.774255	none
0.75, 0.75	-3, -3	0.7	+1	0.727397	none
0.75, 0.75	-3, -3	0.6	+1	0.825511	none
3.60, 0.75, 0.75	+3, -3, -3	0.7	-1.8	1.261795	0.764774
3.60, 0.75, 0.75	+3, -3, -3	0.6	-1.8	1.277780	0.763884

one does not. A test of single readers would miss this failure. Narrower readers preserve the sign against non-valid inputs alone, while a near-resonant valid reader revises an opposing answer despite two conflicting inputs. These trajectories check the mechanism; Section S3.4 proves the uniform bound at $w = 0.6$. Two integrators, Radau and DOP853, agree to below 3.9×10^{-10} at 1001 saved times.

S7.7 Learning misses the sensing interval from both sides (FPL1)

FPL1 asks whether learning from the all-low preparation reaches the functional window and stays in it. It contains sixteen full-network histories at $w = 0.6$ and $\Lambda = 5$, eight with admission and eight with use. The gain $\Lambda = 5$, half the reference strength, is the centre of the writer strengths (2.5, 5, 10) whose numerical accuracy had been established at $w = 0.6$; it was not selected for function. Each history ran for 1250 decision cycles without a state reset at $\Delta t = 0.0005$. Four further histories at $\Delta t = 0.00025$ check the step. FPL1 is a development study: its analyses were made after the histories were inspected.

Exact row-specific sensing intervals. For one row, let x_v be the valid record and d_1, d_2 the two competing records. The selection rule, article Methods Eq. (11), with settled fields gives the valid-source probability p_v exactly at each saved record vector. For the actual d_1, d_2 , the equation $p_v = 0.9$ has a lower and an upper

root, so the sensing interval moves with the competing records and is not a fixed rectangle. At cycle 750, of the 32 valid-link row evaluations, 18 lie below their exact interval, six lie above it and eight lie inside; every admission network has at least one under-developed valid link (article Fig. 4d). At cycle 1250 the counts are five below, seven above and 20 inside. Only U_F02 is in the functional window at the deadline. It first fails at cycle 755, and its longest uninterrupted run of passing sensing endpoints during operation is 46 cycles. No history reaches the functional window by cycle 750 and stays there for the next 500 cycles. Learning therefore fails in two ways: slow acquisition, and over-detuning of a record that stays on the formed branch.

S7.8 A correct outcome causes the exit (FRCP1)

FRCP1 tests whether one positive outcome pulse causes the later sensing failure. Four FPL1 events were selected after inspection: PA1 (A_F04, cycle 869), PA2 (A_F05, cycle 889), PU1 (U_F02, cycle 758) and PU2 (U_F05, cycle 857). The pulse multipliers, comparisons and decision band were fixed before the intervention runs. These are selected events, not independent trials of reliability.

Design. Each event restarts from its complete cycle-start checkpoint: fields, records, memories, traces, baselines, task state and random-generator states. Three arms replay the same first 14 units and then multiply the common six-link writer by $M \in \{1, \frac{1}{2}, 0\}$ during the two-unit pulse. Score, guess status, correction drive, source choices, timing and noise are unchanged, and no state is reset. Two checks protect the comparison: the $M = 1$ arm reproduces the original event, including the exact six end-record bit patterns, and all arms agree in the complete pre-pulse state. All four events pass both checks. The static observation applies the settled-field selection rule to each arm's end records. The live observation evolves a separate two-unit sensing step of the same row from its actual end state, with moving fields and records and matched noise. It was run for $M = 1$ and $M = 0$ only. Each value is the selection probability for one realised field. A replay passes at $p_v \geq 0.9005$, fails at $p_v \leq 0.8995$, and is unresolved in between. All arms use $\Delta t = 0.0005$; no step refinement was run.

Results. Removing the full pulse prevents the failure in all four events, in both observations (Table S9; article Fig. 4e). The relevant record stays above 3.6460668 at every accepted step of all twelve replayed cycles and never crosses x_U , so the failure does not erase the formed record. The pulse moves both the valid record and, through the competitors, the sensing boundary; the largest full-pulse against pulse-off competitor difference in the targeted row is 0.01656, in PA1. The replay shows that the selected pulse causes the exit under these conditions. It does not establish that a positive pulse causes failure from an arbitrary state, nor does it demonstrate optical recovery or repair.

Table S9: FRCP1 sensing probabilities. Static values use each arm's end records; live values use its own two-unit sensing step. Native replay and pre-pulse identity pass in every case. Superscript u: unresolved band. No live $M = \frac{1}{2}$ step was run.

Case	Static p_v			Live p_v	
	$M = 1$	$M = \frac{1}{2}$	$M = 0$	$M = 1$	$M = 0$
PA1	0.895059	0.900313 ^u	0.904681	0.894505	0.904203
PA2	0.892162	0.898027	0.902916	0.893433	0.903737
PU1	0.883100	0.897722	0.907650	0.884537	0.908238
PU2	0.849093	0.884067	0.904127	0.852228	0.904981

S7.9 Sensing lifetimes under continued learning

This section supports the article's statement that learned sensing is lost within at most 151 cycles under the original writer at $w = 0.6$, mostly through the upper boundary. FRCP1 shows that a rewarded pulse can cause an upper exit; this study measures how soon exits occur when the future randomness changes. Using the unchanged original simulation code at $w = 0.6$, $\Lambda = 5$, we continued 14 saved learned states with 128 fresh random streams each (1792 trajectories). Eleven states are the first sensing-capable endpoints of the FPL1 histories that reach that set (four admission, seven use); three further sensing-capable checkpoints are added, and checkpoints that already fail sensing are excluded. Each continuation keeps the complete state (populations, cavity and reader fields, memories, traces, baselines and task state); only the future random streams are new, and nothing is relaxed or reset. The observable is the settled-field selection probability on

both rows, evaluated from article Methods Eq. (11) with the actual competing records; failure is a probability below $0.9 - 10^{-9}$, recorded at cycle ends and also on the integration mesh. The horizon is 500 cycles at $\Delta t = 0.0005$; protection and revision are not monitored, so these are sensing lifetimes. Every continuation exits by cycle 151 at cycle ends and by cycle 125 on the mesh. Conditional median cycle-end lifetimes range from 1 to 39 cycles across the 14 states. Of the 1792 cycle-end exits, 1556 are upper, 227 lower and 9 mixed, so upper loss accounts for 86.8%. This does not contradict the readout bounds, which state what a set of records can do, not whether later outcome pulses keep it in the sensing interval. The 128 futures of a state are not 128 independently learned networks, and zero survivors in 500 cycles is a finite-sample result conditional on these 14 states. Four states with 16 continuations each were repeated at $\Delta t = 0.00025$ with coupled noise increments and matched task randomness: all 64 pairs have the same first-exit cycle and boundary under both monitors, with final-record differences below 4×10^{-7} . The per-state table will be deposited (Section S9).

S7.10 The self-limiting writer in the network (FML1)

FML1 tests whether the writers of Section S6 form useful records, keep both information functions working during continued learning, and replace an obsolete relation. It is a development study at $w = 0.6$ with matched controls, run in a separately written implementation of the published equations, not in the original simulation code. It locates an operating region and tests the design conditions; it does not estimate reliability.

Implementation and validation. The simulator implements article Methods equations (6)–(11), the three traces, the writer, the decision cycle of Section S1.3 and the noise scheme of Section S1.5. It uses fourth-order Runge–Kutta drift steps at $\Delta t = 0.002$, noise increments after each step and seeded random streams keyed by network, cycle and purpose, so all arms and both trace rules of a network use identical external randomness. It reproduces the seven frozen-record reader values of Table S8 to six decimals, and the pulse-time traces of Sections S2.2 and S2.3 to within 1%. At the reference setting ($w = 1$, $\Lambda = 10$) both valid records form in 8/8 admission and 8/8 use networks, with 0/32 and 1/32 persistent non-valid records, and the codes agree event by event on writer increments, selection, assisted-reading rates, the trace at writing and answer revision. Acquisition is faster in the separate implementation at its coarser step, $\Delta t = 0.002$, than in the original code; refinement to $\Delta t = 0.0005$ reduces that difference to sampled formation-time distributions that are consistent between codes; pathwise equivalence remains unestablished. All FML1 counts and the design screen’s event statistics (Section S6) were obtained at $\Delta t = 0.002$; they are development results and are not pooled with Section S7.11. At the FPL1 setting the original writer reproduces the two-sided failure: pooled over both rules, 14 valid links lie below, 6 above and 12 inside their sensing intervals at cycle 750 (FPL1: 18, 6, 8).

Protocol and monitor. Eight networks per arm and trace rule start from the all-low preparation, with the six source assignments cycled over networks. The arms and the monitor were fixed after a four-network admission scan of the write-line parameters, which will be deposited as development data. Each history has 2500 decision cycles with no reset: acquisition to cycle 750, operation to 1250, exchange of need E’s valid source with the unrelated source at 1250, adaptation to 2000 and operation to 2500; all other parameters are those of the reference protocol. The arms reported here are O, the original writer at $\Lambda = 5$; W and S, the original writer at $\Lambda = 1.35$ and 8.75 (the line’s strength at $x = 3.6$ and its peak); P, the write line ($\Lambda_- = 5$, $\Lambda_+ = 8.75$, $\delta_w = 1.25$, $w_w = 1$); PC, the line with the capacity ($x_a = 3.30$, $x_c = 3.60$); G, the line’s gains without its shape ($a_w \equiv 1$); and B, the line on both branches ($\Lambda_{\pm} = 8.75 a_w(x)$). Two capacity-alone arms and four neighbouring line settings were also run; their results, which follow the same pattern, are in the deposit. Every cycle is scored on both rows. Sensing uses the settled-field probability of each row’s valid source every 0.5 time units and the acting row’s actual selection probability from the cavity fields. Protection requires every unassisted reading of a non-valid source that starts with $|\operatorname{Re} z_j| \geq 0.5$ to leave the sign of $\operatorname{Re} z_j$ unchanged with margin 0.5; revision requires every unassisted valid reading to end with the sign of its source value and margin 0.5. A cycle is functional when every test that occurs in it passes at $p_* = 0.9$; a network is accepted when the five cycles before a deadline are functional; strict retention allows no failed cycle in the operating interval. Counts in parentheses use the FRCP1 band, in which values in (0.8995, 0.9005) are unresolved.

Table S10: FML1 acquisition and maintenance, eight networks per cell. Formed: both valid records have crossed x_U by cycle 750. Median formation: median first upward crossing of x_U over the 16 valid links. Accepted: functional in cycles 745–749. Strict: accepted networks with no failed cycle in 750–1249 (values with the FRCP1 band in parentheses). Occupancy, failed-cycle totals and the readout-test counts are in the deposit. Arm letters are those of the FML1 registration.

Arm	Rule	Formed	Median formation (need-cycles)	Accepted	Strict
O, original ($\Lambda = 5$)	A	1	375 (cap)	1	0 (0)
	U	6	130	2	0 (0)
W, uniform weak ($\Lambda = 1.35$)	A	0	375 (cap)	0	–
	U	0	375 (cap)	0	–
S, uniform strong ($\Lambda = 8.75$)	A	7	98	0	–
	U	8	67	2	0 (0)
P, line	A	8	85	8	7 (7)
	U	8	62	8	3 (3)
PC, line with capacity	A	8	85	8	8 (8)
	U	8	62	8	6 (6)
G, same gains, no shape	A	8	69	0	–
	U	8	55	2	0 (0)
B, line on both branches	A	7	107	5	5 (5)
	U	8	74	8	5 (5)

Acquisition and maintenance. The controls behave as the design conditions predict (Table S10; the trajectory figures are in the deposit). The original writer forms slowly and, once formed, sits at the sensing boundary, so that its accepted networks are functional for only a minority of operating cycles. The weak writer never forms a record. The strong writer forms fast, but its valid records settle at 3.9–4.0, above x_+ , and no admission network is functional at the deadline. The line forms both records in all 16 networks and all are functional at the deadline; under admission 7/8 pass the monitor in every operating cycle, the eighth failing in five cycles with settled sensing between 0.8987 and 0.8996. The line with the capacity (PC) forms and is accepted in 8/8 under both rules; under admission all eight pass every operating cycle with no protection or revision failure, and operating valid records span 2.72–3.64. Under use, 6/8 are uninterrupted, both interruptions following downward excursions after large failure pulses (a failed reversal reading at $\Lambda_- = 5$ moves Ψ by about 0.93); no PC network has an upward excursion. The unshaped writer with the line’s gains (G) acquires as fast as the line, but its admission valid records lie above x_+ in 94% of operating cycles and no admission network is functional at the deadline: gain asymmetry alone reproduces the strong writer’s failure. The line on both branches (B) acquires more slowly and maintains.

Replacement. Under the line, with or without the capacity, every obsolete admission record is erased (median 137 cycles) and no false record forms in its place; the new valid record forms a median of about 500 cycles after reassignment, because the new source must be found again by exploration from an all-low row. Under PC, need E is functional at the adaptation deadline in 5/8 admission networks (8/8 under use). The line on both branches (B) erases none of its eight obsolete admission records in the 1250 cycles after reassignment, although under use it erases 8/8 because the use rule’s large failure pulses overcome the attenuated negative branch. The unshaped writer G erases its obsolete admission record slowly (median 749 cycles) and forms the new record in 1/8. These admission results motivate shaping only the positive branch while leaving the negative branch unattenuated. The per-arm replacement table is in the deposit.

Numerical status and limits. Networks 0 and 1 of each rule under P and PC were repeated at $\Delta t = 0.001$. For PC, all four networks are uninterrupted at both steps, and their replacement classifications agree. For P, two of the four strict-retention labels change with the step. The line alone therefore has network-level numerical uncertainty in its strict counts, although occupancy, acceptance and arm contrasts are unchanged. Each cell has eight networks, so exact intervals are wide (7/8: 0.47–0.996). The functional tests sample sensing on a half-unit mesh and test only the readings that occurred; they do not apply the analytical certificates of Sections S3.3 and S3.4. Reader width, trace lifetime, pulse duration, margin and assisted-reading probability were fixed. No capacity other than $x_c = 3.60$ was tested.

S7.11 Six-network validation in the original simulation code (S3VAL)

S3VAL repeats the writer comparison of Section S7.10 on six networks in the original simulation code. The readout tests were registered before the runs. A one-network check in the same code (archive identifier LMR2, to be deposited) preceded it. It supports article Fig. 5 and the article’s statements on acquisition, maintenance, the capacity and replacement.

Design. The code runs at $\Delta t = 0.0005$ and $w = 0.6$ with the admission trace and the schedule of Section S7.10: acquisition to cycle 750, operation to 1250, reassignment of the changed need’s valid source to another source at cycle 1250, adaptation to 2000 and operation to 2500, with no reset. The six S3VAL networks N1–N6 cover the six assignments of the changed need, so they are not replicas of one assignment (Table S12). The five writer laws are the line with the capacity (PC), the line alone (P), the capacity alone at $\Lambda_+ = 8.75$, the line’s gains without its shape ($\Lambda_+ = 8.75$, $\Lambda_- = 5$, $f_w \equiv 1$) and the original writer ($\Lambda = 5$); only the positive branch differs between them. P and PC share their random streams. Each network is scored on six items: acquisition (acceptance over cycles 745–749, an assisted-reading panel at 749 and a full panel at 750); operation (every cycle 750–1249 and scheduled panels); timely replacement (acceptance over 1995–1999 and, at the deadline, sensing of the new source, revision by it and rejection of patterns containing the obsolete source); operation after reassignment (every cycle 2000–2499); the unchanged need (every cycle 1250–2499); and no persistent false record. Sensing passes at $p_v \geq 0.9005$ and fails at $p_v \leq 0.8995$; protection and revision pass at margin ≥ 0.52 and fail at ≤ 0.48 ; a formed record requires $x_{\min} > 2.001$. A panel is a set of readout tests on copies of the network state at one cycle. Acceptance requires the settled sensing test and readings in the five cycles before a deadline to pass at $p_* = 0.9$ and margin 0.5. Scheduled panels add tests between deadlines. A cycle without an eligible reading contributes no protection or revision test; it records coverage only. Each need must nevertheless have at least one executed sensing, protection and revision test in every stage. Otherwise, the item is untested and does not pass. Table S11 specifies the panel schedule, the tests in a panel and the thresholds, as registered before the runs (the monitor specification and the S3VAL registration in the deposit).

Table S11: Scoring protocol of the six-network validation, inherited unchanged by the fresh comparison of Section S7.13. Cycle onsets are counted from zero.

Element	Specification
Deadline panels	A full panel at onsets 750, 1250, 2000 and 2500, paired with a forced-copy-only panel at 749, 1249, 1999 and 2499 so that both rows receive forced readings. Acceptance for the acquisition and replacement items uses the settled sensing test and the eligible readings of the five cycles before the deadline (745–749 and 1995–1999).
Scheduled panels	Full panels during operation at onsets $c = 25m + (m \bmod 2)$, $m = 30, \dots, 99$, excluding the deadlines (67 onsets); during acquisition at $c = 50i + (i \bmod 2)$, $i = 1, \dots, 14$, which are descriptive and do not enter the acquisition item. An empty eligible set at a scheduled panel is recorded as coverage and neither passes nor fails.
Inactive-row sensing	A two-unit clone at the sensing onset of every cycle 750–2499, with the probe of the inactive need on, on its own noise stream; this is the only full-field sensing value for the inactive row.
Preparation	Every test runs on a copy of the complete running state (records, cavity and reader fields, memories, traces, baselines and task state) taken at the onset; nothing is relaxed or reset. Ordinary clones fork at the reading onset, after sensing and before the source draw, with no forcing, pulse or feedback.
Ordinary patterns	All 26 non-empty subsets of the three sources with source values ± 3 ($\sum_k \binom{3}{k} 2^k$), both rows scored in every clone: per row, 8 patterns contain no valid source (protection) and 18 contain the valid source (revision), under the assignment in force (both assignments at onset 1250).
Forced copies	Six copies for the acting row, one per source k with value ± 3 and that reader forced to resonance; scored as installed-reading diagnostics (copy or corruption; valid or non-valid datum), never as ordinary operation. The inactive row in a forced copy is scored as ordinary.
Thresholds and bands	Sensing: nominal $p_* = 0.9$ with band ± 0.0005 (pass at ≥ 0.9005 , fail at ≤ 0.8995 , undecided between). Memory margin: nominal $m = 0.5$ with band ± 0.02 (pass at ≥ 0.52 , fail at ≤ 0.48). Revision: the sign of the presented source value times $\text{Re } z_j$ at the end of the reading, against the margin. Protection: prerequisite $ \text{Re } z_j \geq 0.52$ at the reading onset, then the minimum over integration-step endpoints of the signed margin stays at or above 0.52. Every executed test, including those in the acceptance windows, is scored with these bands; an undecided result does not pass. Formed record: $x_{\min} > 2.001$. Persistent false record: a non-valid link with end-of-cycle $x > x_U$ for more than 20 consecutive cycles, the obsolete link exempt in 1250–1999 with its counter restarting at 2000.

Cycles are numbered from 0, and a quantity assigned to cycle c is evaluated on the state at the end of that cycle; cycles 2000–2499 are therefore the 500 cycles of operation after the change. For each link outside the current valid set, the false-record item counts consecutive cycles ending with $x > x_U$ and fails at 21. The obsolete link is exempt during adaptation (cycles 1250–1999), and its count restarts at cycle 2000. The new

link is valid from cycle 1250. Under PC both failures of this item are the obsolete link of the changed need, whose first uninterrupted run of cycle ends above x_U , starting when the count restarts at cycle 2000, lasted 66 (N5) and 187 (N6) cycles; no other link failed it. A network completes the sequence when all six items pass. Two separately written scorers agree on all 180 item labels and 30 conjunctions.

Table S12: S3VAL item outcomes for the line with the capacity (PC) and the line alone (P). Items in order: functional acquisition, operation, unchanged need, timely replacement, operation after the change, no persistent false record. +, pass; −, fail. The assignment gives the changed need and its obsolete and new valid source. The last column is the first failure under PC, by cycle.

Network	Assignment	PC	P	First PC failure
N1	E: 1→3	+ + + − − +	− + + − − +	replacement at 1995
N2	I: 2→3	+ + + + + +	+ + + + + +	−
N3	E: 2→1	+ + + − − +	+ + + − − +	replacement at 1995
N4	I: 3→1	+ + + + + +	+ + + + + +	−
N5	E: 3→2	+ + + − − −	+ + − − − +	replacement at 1995
N6	I: 1→2	+ + + − − −	+ + + − − −	replacement at 1995
Networks passing each item		6 6 6 2 2 4	5 6 5 2 2 5	

For the controls, the capacity alone passes acquisition and operation in 6 of 6 networks but completes the sequence in none; the line’s gains without its shape and the original writer fail acquisition in 6 and 5 of 6 networks and complete none. The line alone fails acquisition in N1 by a margin (settled sensing 0.8984 and 0.8987 against 0.9 in the acceptance cycles) and the unchanged need in N5 through one inactive-row sensing test (0.89919); its complete sequences are the same two as for PC. The registered complete-sequence contrasts are not significant at $n = 6$ (exact McNemar $b = 2$, $c = 0$, $P = 0.5$ against each of the three failing controls). The registration paired the arms by network, that is, by assignment and preparation; only P and PC also share random streams. The corresponding unpaired comparison, two of six against none (Fisher’s exact test, two-sided $P = 0.45$), is also not significant. The article’s functional-acquisition contrast, six of six under PC against none under the line’s gains without its shape, was not registered. These arms do not share random streams, so the comparison uses an unpaired Fisher’s exact test (two-sided $P = 2/924 = 0.002$). The exact 95% Clopper–Pearson interval for six of six is $[0.54, 1.00]$. Against the original writer, which acquires in one of six, $P = 0.015$. This study tests whether the complete sequence can occur at one parameter point; it does not estimate reliability.

The capacity suppresses upper excursions. Upper-excursion time is summed over all six links in cycles 0–2499 using per-step counters, in time units. A counter starts when a record exceeds the chosen level and stops when it falls 0.01 below it. Article Fig. 5b abbreviates this as time above the level:

	N1	N2	N3	N4	N5	N6	Total
P, $x > 3.60$	4417.1	8561.6	6251.5	11428.6	8123.7	5362.4	44144.95
PC, $x > 3.60$	57.2	176.1	225.8	255.2	225.3	83.9	1023.45
P, $x > 3.70$	92.6	73.8	231.9	365.9	209.8	3.8	977.80
PC, $x > 3.70$	0	0	0	0	0	0	0

The reduction above 3.60 is between 28-fold (N3) and 77-fold (N1) in every network, 43-fold in total. The family maxima of x are 3.711–3.754 for P and 3.628–3.660 for PC: positive writing alone cannot carry a record across x_c (Section S6), but the full noisy dynamics can. Time above a population level is not time spent failing a readout test. Because $\kappa_{\text{cap}} = 1$ for $x \leq x_a = 3.30$, P and PC have the same positive writing below x_a and the same negative writing everywhere, which is consistent with their identical replacement outcomes.

The obsolete record keeps control of selection (analysis after the runs). Table S13 follows the changed need from the first crossing of the obsolete record below x_U ($x_{\min} \leq 2$) to the deadline, under PC.

Table S13: Replacement under PC. Selection probabilities are means over the changed need’s decisions from the first crossing to cycle 1999; “obsolete above x_U ” counts the cycles in that interval that end with the obsolete record above 2. Exposure counts the 375 decisions of the changed need during adaptation (cycles 1250–1999): selections of the new source, and writer-enabled assisted readings of it (assisted reading, non-guess action), with the number from the uniform-exploration branch in parentheses.

Network	Result	First crossing	$\bar{p}_{\text{obsolete}}$	$\bar{p}_{\text{new}}$	Obsolete above x_U	New selected	Assisted, writing
N1	fail	1632	0.518	0.290	43 of 368	65	25 (7)
N2	pass	1313	0.210	0.668	0 of 687	232	82 (6)
N3	fail	1460	0.592	0.281	188 of 540	85	32 (9)
N4	pass	1439	0.340	0.565	57 of 561	160	65 (4)
N5	fail	1458	0.879	0.065	379 of 542	25	13 (6)
N6	fail	1405	0.919	0.041	589 of 595	10	5 (5)

In N5 and N6 the obsolete source is the most probable choice at every decision of the changed need after the first crossing. Two effects contribute. First, the obsolete record is not removed: in N6 the first crossing is a brief dip, the record is below 2 at only 3 of the 298 later decisions and returns to about 3. Second, a partly decayed record keeps its advantage: in N5 the record is below 2 at 76 decisions (median 1.82) and the obsolete source is still top-ranked at all 76, because the sensing intensity varies continuously with x and does not switch at x_U . With the obsolete record at $x = 2$, its competitors at x_L and settled fields, the selection rule gives the new source $p = 0.057$. The new source is therefore read rarely, and its record receives few writes: in N6 all 5 writer-enabled assisted readings during adaptation came from the uniform branch, and the new record never ends a cycle above 2 before cycle 2500. In N1, N3 and N5 the new record forms after the deadline (first end-of-cycle crossing at cycles 2256, 2068 and 2444), once the obsolete record has fallen; a later formation cannot satisfy a deadline item. Access is not the only factor. From the cycle after the first crossing to cycle 1999, net writing on the obsolete link was downward in the coordinate Ψ of Eq. (S4.1): -5.90 in N5 and -15.64 in N6. Yet the records rose from $x = 1.87$ to 2.60 and from 1.95 to 2.86 ($\Delta\Psi = +1.05$ and $+1.30$). Passive filling and restoration, with noise also contributing in N6, outweighed net negative writing. In N3 the new link received substantial positive writing that physical restoration largely opposed, and realised noise was unfavourable in N5. The selection and exposure measures were computed from the raw histories after the runs. The reset intervention in Section S7.12 shows that the obsolete record blocks replacement in N5 and N6. Because the reset lowers both readouts, it does not isolate the effect of access.

Scope. The study uses six networks at one parameter point and one step size. In the later comparison of three writer laws, halving the step left every label for these networks unchanged (Section S9.1). The control arms are reported from their scores, not rechecked against raw trajectories. The outcome signal Q is supplied by the evaluator, as in all studies of this paper.

The first obsolete input. A validity change is unavailable to the material until its inputs or outcomes reveal it, so no local rule can both admit the next valid source value and reject the first newly obsolete one; identical local histories cannot support reliably opposite decisions. The replacement tests therefore measure learning from later outcomes.

S7.12 Releasing the obsolete record (RLR1)

Section S7.11 leaves open whether the obsolete record itself blocks replacement in N5 and N6 and whether a local rule can remove it. We address these questions with development tests in the original simulation code, on the S3VAL networks and parameters ($w = 0.6$, $\Delta t = 0.0005$, admission trace, the PC writer otherwise unchanged). All six S3VAL networks were already exposed when these studies were designed, so they are development results, not validation.

The extension: an adverse-outcome susceptibility. Each link receives one further state μ_{jk} , signed and unbounded, with no noise, set to zero at the start of the unformed preparation:

$$\dot{\mu}_{jk} = \sigma\epsilon_{jk}[\Theta(-Q) - r_\mu\Theta(Q)] - \frac{\mu_{jk}}{T_\mu}, \quad \Pi_{jk} = \sigma\epsilon_{jk}H(x_{jk})\{\Lambda_+a_w\kappa_{\text{cap}}[Q]_+ - \Lambda_-(1 + \eta_\mu[\mu_{jk}]_+)[-Q]_+\}, \quad (\text{S7.1})$$

with Θ the unit step and all other quantities as in article equation (4). The susceptibility charges when admitted readings through the link are followed by wrong answers and discharges after correct answers. It

multiplies only the negative branch, adding depletion without positive credit. Provided $0 < \beta + \vartheta < 1$, the sign of Q identifies success and failure. The mean outcome-driven input per admitted reading is then proportional to $(1 - p) - r_\mu p$, where p is the link’s trace-weighted success. This input changes sign at $p_0 = 1/(1 + r_\mu)$. The total drift also includes the relaxation $-\mu/T_\mu$. With $r_\mu = 1/3$, the threshold is $p_0 = 0.75$, independently of β within that range. The baseline update confines β to $[0, 1]$. Because β estimates fair-guess success, it stays near 0.5, well below $1 - \vartheta = 0.8$. Over all cycles of the 180 VAL1 histories, it lies between 0.33 and 0.68 (mean 0.497, standard deviation 0.047, against 0.050 for the stationary update). It is identical under the three writers because only designated guesses update it. The law uses the link’s own trace, population and susceptibility, together with the broadcast pulse flag and the sign of the broadcast outcome signal already used by Eq. (4). It receives no validity label, source identity or reassignment time. For links with $\mu_{jk} \leq 0$ it equals the PC writer, and with $\eta_\mu = 0$ it reproduces the PC histories exactly on the physical fields. The parameters were fixed before any run with this law. The choice $r_\mu = 1/3$ places the break-even success between chance (0.5) and valid-link success (0.85–0.94). The lifetime $T_\mu = 2200$ time units (100 cycles) was chosen with knowledge of the reassignment schedule. The gain $\eta_\mu = 1.3$ came from an empirical screen using the measured negative-writing balances of the obsolete links in N5 and N6. A first candidate, negative writing enhanced in a band of populations, was rejected before any run: initial acquisition dwells for 200–300 cycles inside the only admissible band, and the obsolete record of N6 lies above it, so a population-dependent law cannot treat the two differently.

Premise check on the saved histories. Before any run, the susceptibility was computed along the saved PC histories with the populations held fixed (read-only, two independent implementations agreeing to 4×10^{-16}). After the first crossing of x_U , trace-weighted success was 0.558 for the obsolete link in N5 and 0.553 in N6, both below p_0 . In the 18 windows of formed valid links, it was 0.838–0.940, and $\mu < 0$ at every cycle start, so the extension added no negative writing. In forming links it was 0.796–0.876, and the added dose was at most 5.6% of baseline.

Reset of the obsolete record (O1). The PC histories of N5 and N6 were restarted from complete saved states at cycle 1626, a boundary between saved segments. Two arms used the same random streams: unchanged continuation (C) and reset of the obsolete record alone to x_L , with the writer unchanged (O1). The reset is a diagnostic, not a learning rule. The following table gives the changed need over its 187 decisions in cycles 1626–1999 (two independent reconstructions agree on every entry):

	N5		N6	
	C	O1	C	O1
Obsolete record at cycle 1626	1.837	0.586	2.667	0.586
Obsolete record at cycle 2000	2.599	0.316	2.865	0.732
Decisions with new sensing intensity higher	0	1.000	0	0.989
Mean selection probability, obsolete/new	0.88/0.07	0.08/0.83	0.92/0.04	0.10/0.81
Writer-enabled assisted readings of the new source	6	53	2	75
First cycle with the new record above x_U	none	1736	none	1717
New record at cycle 2000	0.379	3.284	0.817	3.568
Minimum of the unchanged valid record	2.991	3.062	2.984	3.030
Replacement test	fail	pass	fail	pass

Resetting the obsolete record alone was enough for the unchanged writer to acquire the new source and pass the replacement test in both networks, and the unchanged need lost no sensing test. The reset lowers both readouts of the record, so it establishes that the record blocks replacement, not that the sensing readout alone does.

From the unformed preparation (six networks). Eq. (S7.1) was then run from cycle 0, with $\mu = 0$, on all six S3VAL networks with their S3VAL random streams, and scored by the two scorers of Section S7.11 (agreement on all 36 labels). The extension passed all six items in five networks, against two under PC, and was never worse on any item; in N1 the new record formed after the deadline under both writers. The persistent false records of PC in N5 and N6 did not occur under the extension. The obsolete link’s susceptibility became positive after reassignment and relaxed towards zero by cycle 2000, and its total negative dose during adaptation was lower than under PC in five networks, consistent with earlier release; formed valid links had $\mu < 0$ at every cycle of operation and of the unchanged need (18 records). One expectation fixed before the runs failed: that the first formation times of the two initially valid links stay within $\pm 10\%$ of PC. Three of twelve links fell outside it (one later, two earlier); all still passed functional acquisition. Under the pre-run decision rule, this failure meant that the extension was not recommended for fresh validation. A fresh

three-arm validation was nevertheless commissioned as a separate decision, while retaining the development verdict.

Scope. The reset uses two networks; the study from the unformed preparation uses six. All were exposed during design, at one parameter point and one step size. One parameter of the extension was screened on N5 and N6. These results show that the obsolete record blocks replacement in the two severe cases and that a local susceptibility, without validity labels, can release it. Section S7.13 tests whether the extension improves on PC or a uniformly stronger negative branch in histories not used in its design. The susceptibility needs its own physical channel, for example a slow, outcome-driven second population of the same link.

S7.13 Fresh three-arm comparison (VAL1)

This section reports the confirmatory test that Section S7.12 called for: the extension of Eq. (S7.1), with its parameters unchanged, against the self-limiting writer and against a uniformly strengthened negative branch, on 60 networks not used in its design. The cohort, endpoint, control, tests and decision rule were fixed and recorded in a dated design record before any new history was run; that record will accompany the public deposit.

Design. Sixty new networks, ten for each of the six reassignment classes in Section S7.11 (initial valid sources and changed row), run from the unformed preparation in the original code at $\Delta t = 0.0005$. They follow the S3VAL schedule: acquisition deadline at cycle 750, reassignment at 1250, replacement deadline at 2000, and end at 2500. Each network is run under three writers with matched random streams: the self-limiting writer (PC), article equation (4); the extension (F1), Eq. (S7.1), with $r_\mu = 1/3$, $T_\mu = 2200$, $\eta_\mu = 1.3$ and $\mu = 0$ at preparation; and a uniform control (U), which keeps PC's positive branch and multiplies its negative gain by a constant. The development histories never enter the denominator.

The uniform control. The uniform control tests whether the extension's improvement can be explained by stronger negative writing alone. The control gain is calibrated on the six saved PC development histories of Section S7.11, using a read-only calculation before any new run. Let $dI^- = \Lambda_- \sigma \epsilon[-Q]_+ dt$ be the baseline negative dose in the coordinate Ψ of Eq. (S4.1). The susceptibility was filtered passively along every link of all six histories over cycles 0–2499, giving

$$\eta_U = \frac{\sum \eta_\mu [\mu]_+ dI^-}{\sum dI^-} = 0.2284, \quad \Lambda_{-,U} = \Lambda_- (1 + \eta_U) = 6.142, \quad (\text{S7.2})$$

with the sums over all links and histories (baseline dose 603.53, added dose 137.85). On the development paths, the control therefore adds the same total negative dose as the extension, spread uniformly over links, populations and time. The live paths subsequently diverge, so this identity calibrates the control but does not predict equal doses in new networks. It supports no energy or dose-efficiency claim.

Endpoint and tests. The single primary endpoint is the complete sequence: $Y = 1$ if and only if all six S3VAL criteria pass in the same history. They are functional acquisition, operation, replacement, operation after reassignment, the unchanged need and no persistent false record. The two S3VAL scorers agreed on all 1,080 component labels and 180 conjunctions. A history that the scorer could not resolve numerically counts as $Y = 0$ and is reported separately; it cannot be reclassified. The two primary contrasts are F1 against PC and F1 against U. They use the two-sided exact McNemar test on discordant pairs, with Holm adjustment over exactly these two tests at familywise level 0.05. Superiority also requires a positive paired difference. U against PC is a secondary contrast, reported with its unadjusted P value only. The cohort size was fixed at 60 by a power calculation for a large effect. The timing expectation that failed in development (first formation times within $\pm 10\%$ of PC on the two initially valid links) is reported here as a secondary diagnostic; its development failure stands.

Primary result. All 180 histories completed with no error or missing family.

Writer	Complete sequence	Rate	Explicit fail	Unresolved (counted as fail)
PC	35/60	0.583	23	2
F1	47/60	0.783	11	2
U ($\Lambda_{-,U} = 6.142$)	41/60	0.683	17	2

Paired contrasts (first writer only / second writer only; both pass; neither):

Contrast	First only	Second only	Both	Neither	Difference	Exact P	Holm P
F1 against PC	13	1	34	12	+0.200	15/8192 = 0.00183	0.00366, rejected
F1 against U	7	1	40	12	+0.100	9/128 = 0.0703	0.0703, not rejected
U against PC (secondary)	10	4	31	15	+0.100	1471/8192 = 0.180	not in the family

Decision. Under the rule fixed before the runs, the outcome is the second of four: F1 improves on PC, and superiority to U is not established. The data support an improvement of 20 percentage points in the complete sequence over the self-limiting writer at the tested parameter point, task, step and horizon. They do not establish an advantage of history dependence over this fixed gain calibrated to the development dose. Seven of eight discordant pairs favour F1, and the point estimate is +0.10, but the registered test does not reject. Non-rejection does not establish equivalence. Any statement about U applies to this calibrated gain, not to every constant gain. The cohort may not be enlarged to pursue the second contrast.

Components. Passes of each criterion, of 60 (PC/F1/U): acquisition 56/59/51; operation 52/54/47; replacement 43/57/54; operation after the change 42/56/53; unchanged need 58/58/57; no persistent false record 52/59/59. F1’s gain over PC lies in replacement, operation after reassignment and avoidance of persistent false records. U shares most of the replacement gain but, relative to PC, loses acquisition and operation in six networks each and gains each criterion in one other network. Those two criteria account for F1’s gain over U. These component differences are descriptive; no component replaces the complete sequence as the endpoint.

Timing diagnostic. First formation times of the two initially valid links were within $\pm 10\%$ of PC in 104 of 120 link-networks for F1 (12 outside, 4 censored); the counts for the other pairs are in the deposit. The other secondary readouts (admissions of the new source, dwell times, erasure times and realised doses) will be deposited; across writers they follow the expected sequence of earlier erasure and more admissions, but within F1 they do not separate successes from failures, and the lower realised dose of F1 is not an efficiency result because U was matched on the development paths only.

Evidence. An independent audit recomputed every count, discordance, exact probability, Holm adjustment and component table from the saved records and found all of them to agree; it did not rerun the scorers on the raw histories. The design lock, the results archive and the audit will be deposited (Section S9).

Scope. The study covers one task, one parameter point and one step, with one halving checked on the development networks only. Its horizon is 2500 cycles, and it uses supplied selection, outcome and assisted-reading signals. No physical channel is demonstrated for the capacity or susceptibility. The control is matched in added dose only on the development histories. The extension completes the sequence more often than the self-limiting writer in networks not used for design. The failure modes of Section S7.11 are reduced but remain (11 explicit failures of 60). The study does not establish that history dependence is necessary, optimal or more efficient than a constant gain. The failed development expectation for formation timing also stands.

S8 Physical implementation

This section separates demonstrated device functions from the connections still needed to implement the model. Section S8.1 describes a candidate link and compares its requirements with experiments; Section S8.2 maps device quantities onto the equations; and Section S8.3 sets out the decisive tests. No learning cell of this kind has been built; every requirement below is a target for the proposal.

S8.1 A candidate polariton learning link and its experimental support

The candidate link is a confined polariton structure, such as a micropillar or a pair of coupled micropillars etched from an InGaAs/GaAs quantum-well microcavity at 5–15 K, in which two weak optical modes overlap one long-lived exciton reservoir. Table S14 lists, for each state of the model, what the device must provide, the proposed carrier with its experimental precedent, and the step still missing. The quoted values come from

different samples and operating points; combining them would be a proposed integration, not a measured device.

The first prototype keeps the trace and outcome gate electronic. It tests the model’s central physical claim: one learned material state regulates both which source is read and how strongly it changes memory. Local analogue electronics can supply eligibility and gating for this test. An all-optical trace is possible in principle but remains the least mature component (Section S8.2).

Table S14: Requirements of the candidate link, their proposed carriers with experimental precedents, and the step still missing for this model. Reported values are from the cited experiments; design scales are marked.

Model requirement	Proposed carrier and precedent	Missing step
Relevance record n : one persistent population that shifts both optical readouts and can be raised and lowered by the writer	Long-lived (dark or localised) quantum-well excitons held by the continuous holding drive. A dark-exciton reservoir tracked through the lower-polariton blueshift in an InGaAs/GaAs microcavity [8]: lifetime above 20 ns against a polariton lifetime of order 10 ps, linewidth about 170 μeV , shift of order 40 μeV at the reported loading, strong coupling retained at about 175 pJ per loading pulse and degraded at twice that. A slow reservoir fed by the resonantly driven fluid, $\dot{n}_R = -\gamma_R n_R + \gamma_{\text{in}} \psi ^2$ [9], with a component of about 400 ps	The model’s excursion from x_L to x_H is $2\sqrt{2}$ half-widths, about 240 μeV at this linewidth, six times the reported shift; the shift and the time-scale separation must be met jointly (Section S8.2); the intracavity-loading law of a long-lived species remains to be established
Sensing mode b : a resonantly driven mode whose probe response reports n , with the record stored by optical bistability	The polariton mode of the pillar at the blue detuning of bistability, read through one polarisation or transverse mode. Dark excitons narrow and finally collapse the hysteresis of a resonantly driven polariton state [10] (upper threshold moved by about 150 mW over the reported sweep)	The reservoir does not only translate the resonance; the functional window must be measured as the range of n over which the bistability itself survives
Reader c : a second mode, coupled to the source and the answer memory while the source is selected, whose coupling depends on n	The orthogonal polarisation of the same mode (splitting about 30 μeV , locally up to 100 μeV [11]; spin-resolved bistability under nanosecond excitation [12]), or a second transverse mode or coupled pillar overlapping the same reservoir	Two modes of the same transverse profile have similar excitonic fractions and shifts; $w = 0.6$ then needs $g_r/g_R \simeq 3.3$ or a narrower reader (Section S8.2); cross-talk through the shared reservoir and spin-dependent interactions must be measured
Positive writing: trace- and outcome-gated loading of n proportional to $H(x)$, stopping at a ceiling inside the window	Non-resonant femtosecond loading of the dark reservoir [8] (pulses of order 100 pJ); the resonant write tone of Section S6 (design scale: write power of order ten times the holding power during the pulse)	The write must be set by the link’s own retained trace and the outcome; no state-dependent self-limiting write has been demonstrated
Negative writing: trace- and outcome-gated depletion of n	Carrier-assisted optical depletion of a non-radiative exciton reservoir in GaAs/AlGaAs quantum wells [13]; local exciton ionisation under an electrical gate in a bistable polariton transistor [14]	Shown as control operations in a bare quantum well and a transport device; the response time and graded update law in the strongly coupled link remain to be measured
Assisted reading: temporary tuning of the selected reader to resonance	Dynamic Stark shift of strongly coupled polaritons [15] (approaching 1 meV at 2 mJ cm ⁻² , Rabi doublet unaffected)	Must be applied per link and per reading on the selected reader only
Trace ϵ : a local state that retains the link’s reading activity until the outcome arrives	In a first prototype, a photodiode on a weak reader tap and a leaky integrator with lifetime of order ten memory response times, enabled only while the link’s need is acting (a 10 ns RC trace retains $e^{-0.2}$ over a 2 ns outcome delay, design scale); eligibility traces carried by resistivity drift in phase-change devices [16]	A material carrier whose relaxation into the reservoir is negligible before the outcome (Section S8.2)
Answer memory z : two stable phases, 0 and π , coupled reciprocally to the readers	A phase-bistable condensate or degenerate parametric oscillator (Section S8.2). Spin-bistable condensates stable for seconds and switched by an opposite-spin pulse about 100 times weaker [11], up to 95% circular polarisation at 5 K	A spin-bistable state is a different element from the phase-bistable memory and would need its own reduction of the reader coupling
Susceptibility μ (extension only): a slow signed state integrating adverse outcomes	A local electronic integrator in a prototype	No material carrier identified

Three constraints follow from the table. First, record excursion, survival of strong coupling and time-scale separation must be satisfied together. The required reservoir shift scales with linewidth: narrowing the mode lowers the required density but slows the optical response, reducing κ/γ at fixed reservoir lifetime (Section S8.2). Second, loading and depletion have been demonstrated as control operations in different samples. The model needs both as graded, trace-gated updates in one strongly coupled link. Third, photothermal shifts are a poorer match for this record than an exciton reservoir. The reservoir experiment above distinguished

dark-exciton loading from heating by the sign of the shift; heating moved the resonance in the opposite direction.

S8.2 Mapping to the equations

Time scales. Here κ is the amplitude decay rate in Eq. (S1.3). The intensity linewidth is $\Gamma_{\text{FWHM}} = 2\hbar\kappa$, the unit of x is the half-width $\hbar\kappa$, the amplitude lifetime is $1/\kappa$, and the population lifetime is $1/(2\kappa)$. A reported polariton lifetime must be assigned to the appropriate convention before forming ratios. The reference parameters set relative, not absolute, time scales: in memory response-time units the cavity field relaxes in $1/\kappa = 0.01$, a reading lasts 10, a cycle 22, the reservoir decays in $1/\gamma = 1000$ and the susceptibility in $T_\mu = 2200$. One illustrative mapping: a polariton amplitude lifetime of 10 ps for $1/\kappa$ gives a memory response time of 1 ns, a decision cycle of 22 ns, an acquisition period of 750 cycles of about 17 μs , and a reservoir lifetime of 1 μs . The reference ratio $\kappa/\gamma = 10^5$ is a parameter choice rather than a requirement. Ref. [1] derives how this ratio changes condensate dynamics, from a fast reservoir that follows the field to a slow one that does not. The measured dark-exciton reservoir [8] gives a ratio of about 2×10^3 against a 10 ps amplitude lifetime. The model forms records at $\kappa/\gamma = 2500$, the slowest field setting in the scan (Section S7.5, Table S7), but the complete sequence in Section S7.11 has not been tested there. The required shift and time-scale separation must be met together. Narrowing the mode lowers the required shift but also slows the optical response. For the reference excursion $\Delta E = 2\sqrt{2}\hbar\kappa$ and reservoir lifetime τ_n , the relation is $\kappa/\gamma = \kappa\tau_n = \Delta E \tau_n / (2\sqrt{2}\hbar)$. If the shift were limited to the reported 40 μeV and the lifetime to 22 ns, then $\kappa/\gamma \lesssim 470$, below the smallest tested ratio for formation. This conditional calculation is not an impossibility result. The reported shift is not an established maximum, and different storage parameters or reservoir species change the requirement. Device design should therefore begin with a rerun at the measured ratios.

Reader width. If the cavity shift is $g_R n$ and the reader shift is $g_r(n - n_H)$, then $w = (\kappa_r/\kappa)(g_R/g_r)$. The reference values $\kappa_r/\kappa = 2$ and $w = 1$ require $g_r = 2g_R$. At the same linewidth ratio, the functional design $w = 0.6$ requires $g_r \simeq 3.3g_R$. A reservoir shifts a polariton mode in proportion to its excitonic fraction. Two polarisations of one transverse mode therefore give $g_r \simeq g_R$, in which case $w = 0.6$ requires $\kappa_r/\kappa \simeq 0.6$: the reader must be narrower than the cavity. Two transverse modes or coupled pillars offer more freedom in excitonic fraction and linewidth, but both must be shown to sample the same n . The decay rates κ , κ_r , shift coefficients g_R , g_r and reservoir overlaps must be designed and measured together. The measured width $w_{\text{exp}} = (\kappa_r/\kappa)(g_R/g_r)$ then determines, through Section S3.3, whether protection and revision share a window. The readout peaks are separated by $x_H - \delta = \sqrt{2} - 1 \simeq 0.414$ half-widths in the record coordinate. At the 170 μeV full width above, this corresponds to about 35 μeV of shift. It does not specify the splitting of the bare optical modes, which also depends on their drive frequencies and reservoir couplings.

Which term a control changes. Each physical control changes a particular term in Eqs. (S1.8)–(S1.11). A loading beam or absorption of cavity light adds to $\dot{n}$ (positive Π or the $|b|^2$ term). Carrier-assisted depletion or ionisation removes from $\dot{n}$ (negative Π). A Stark shift or gate field that moves a resonance without changing n changes the detuning (Δ , Δ_r), as in assisted reading. Added loss changes κ or κ_r . Only the first two controls write the record. A temporary transmission change therefore does not demonstrate negative writing, and electrical control of the optical response does not by itself establish a persistent negative update of n .

One scalar coordinate. The model assumes that one scalar population shifts both readouts. In a spin-resolved device, the reservoir's two circular components interact with the polariton field at different strengths, so the channels can shift differently [12]. An overlapping reservoir can also change condensate polarisation, rather than energy alone [11]. The first measurement must therefore test whether both readouts are functions of one inferred n (Section S8.3). If separate spin populations are needed, the link has an additional state variable. The reduction to these equations then holds only where the populations remain proportional.

Field linearity at fixed reservoir. The element equations and the two-tone writing calculation treat the cavity field as linear when the reservoir population is held fixed: the only shift is $g_R n$. Polariton experiments are commonly modelled with an additional direct self-interaction shift, $g_{\text{pp}}|a|^2$, alongside the reservoir term [8], so the field equation would read $\dot{a} = [-\kappa + i(\Delta - g_R n - g_{\text{pp}}|a|^2)]a + f$. Recovering the model therefore also requires $|g_{\text{pp}}||a|^2 \ll \kappa$ throughout holding, probing and writing, with the residual shift small enough for the claimed readout accuracy; the writing stage deserves particular attention because its proposed power is of order ten times the holding power (Section S6). A long reservoir lifetime does not by itself secure this condition: at quiet equilibrium $n = (s_{\text{cap}}/\gamma)|a|^2$, so the ratio of the two shifts is $g_{\text{pp}}\gamma/(g_R s_{\text{cap}})$ and depends on

the capture rate and the interaction coefficients. Where the condition fails, the self-interaction term must be retained when deriving the bistability, the readout curves and the averaged write response. This is a condition on the proposed implementation; it does not alter the shared-record mechanism of the effective model.

Tone phase. The weak probe and read tones must stay in quadrature with the holding drive (Section S1.2). With both tones on, a phase error φ towards the holding drive changes the power by $4\sqrt{A_0 B} \sin \varphi$; staying below the upper fold requires $|\varphi| < 4.6^\circ$ at $\delta = 3$, $B = 0.2$, in the quasistatic limit, so the tones must be phase-locked to the holding drive.

Answer memory. Write $z_j = q_j + ip_j$ and omit the source and correction drives in Eq. (S1.11). Then

$$\dot{q}_j = (1 - q_j^2 - p_j^2)q_j, \quad \dot{p}_j = -(G + q_j^2 + p_j^2)p_j.$$

For $G > 0$, the imaginary part decays and the stable states are $z_j = \pm 1$. For $G > 1$, as at $G = 80$, the linear z_j coefficient is a loss, the z_j^* coefficient supplies a larger phase-sensitive gain, and the cubic term limits the amplitude. This is the normal form of a resonant degenerate optical parametric oscillator above threshold after eliminating the parametric pump [17]. The model thus uses two storage mechanisms: a reservoir population for relevance and opposite optical phases for the answer. A spin-bistable condensate [11] stores distinguishable states, but they are not opposite phases of one field. Replacing the answer memory by such a condensate would require a separate reduction of its reader coupling.

Reader coupling. The terms $-ivI_k z_j$ and $-ivI_k c_{jk}$ exchange amplitude between memory and reader, and the source enters in opposite phase as $-ivI_k(z_j - y_k)$ [18]. This needs compatible field normalisations and a controlled relative phase. With one link, no source and the memory's own terms omitted, $\frac{d}{dt}(|z_j|^2 + |c_{jk}|^2) = -2\kappa_r |c_{jk}|^2$, so the coupling is passive. The term $\gamma\chi |c_{jk}|^2$ assumes that reader absorption fills the same reservoir; its loss must be included in κ_r .

Trace. The reader field decays too quickly to serve as the trace. Its lifetime is $1/\kappa_r$, leaving an amplitude fraction e^{-400} over the two-unit outcome delay, whereas the trace retains $e^{-0.2}$. A material trace could use a second, shorter-lived population n_e filled by reader light; nanosecond free-carrier and bright-exciton components exist [13]. If n_e relaxes into the persistent reservoir, however, reading writes the record before any outcome arrives. This would break the separation between tagging and writing a link. A material trace therefore needs negligible coupling to the reservoir until the outcome gate opens, for example in a separate well or carrier species. This isolation must be measured. The admission gate P_j would require efficient reader absorption only while its row is illuminated; that mechanism is not derived here. The first prototype therefore keeps the trace electronic.

Writer, capacity and susceptibility. The positive branch uses the resonant write tone of Section S6, with pulse power of order ten times the holding power. The capacity requires a separate filling channel with finitely many states. Ideally it would appear as measured saturation of loading, $\Delta n_+(n) \rightarrow 0$ as n approaches a limit; such a channel can be represented by a population-dependent factor only if its occupation follows n rapidly, and otherwise it is an additional dynamical state (Section S6). The susceptibility is a slow signed state that a local integrator can carry; no material carrier is identified.

Scope. The holding drive remains on during retention. At the stated phase and power, the weak tones stay below the switching fold. Their disturbance and the reader's loading must nevertheless be measured in a device. The derived target is record drift below about 0.01 per reading, so that ten unassisted readings move a record near x_H by less than its 0.11 margin to the window's lower edge. No energy or speed advantage over a digital controller is established.

S8.3 Decisive experimental tests

Each test addresses a missing connection, ordered so that a failure informs the next design decision. The first tests the shared-record mechanism without learning.

1. *Shared-record test (one link, no writer).* Prepare different reservoir populations with a calibrated loading beam. At each, measure the sensing mode's probe response, the reader's coupling to a test memory, both linewidths, the Rabi splitting and the hysteresis boundaries. Pass: both responses are reproducible,

single-valued functions of one inferred n over an interval where bistability survives, and both channels yield the same n . This implements the numerical transplant of Section S7.3: a power pulse above the write fold sets a high branch, and a pulse below the erase fold sets a low branch. A chosen record pattern can then be prepared and both readouts tested, with memories, traces, fast fields and inputs matched. Equal records provide the control. Fail: channel histories cannot be described by one common state, showing that the one-population assumption is incomplete.

2. *Bidirectional writing with the readouts preserved.* At several starting populations, apply calibrated loading and depletion pulses and remeasure both spectra to map $\Delta n_{\pm}(n, E_{\pm})$. Then cycle 10^2 – 10^3 paired updates about three set points and remeasure resonances, linewidths, Rabi splitting, hysteresis thresholds and both contrasts. Pass: graded updates in both directions preserve sensing and reading contrast. Saturation of Δn_{+} at high n would supply the capacity, provided the saturating channel follows n rapidly (Section S6).
3. *Read-without-write and locality.* Repeat sensing and reading without an outcome and measure the drift of n per reading, expressed as a change in the normalised record (target below 0.01 cavity half-widths, Section S8.2). In a multi-link sample, write one link and measure the spectral shift of a neighbour. Exciton diffusion over tens of nanoseconds can defeat optical confinement.
4. *Delayed-writing test (three-factor rule).* With the electronic trace in place, inject known reader pulses and measure the trace retained until the outcome delay. Test all four combinations of reader activity present or absent and outcome present or absent. Substantial writing should require both and should be proportional to the retained trace.
5. *Two-source learning (two links, one memory).* Let one source carry the relevant signal and the other an unrelated signal. Deliver a common outcome after each action. Let selection circuitry sample from the measured sensing intensities without storing an estimate of source reliability. Measure valid-source selection probability and each source's influence on the memory against the two reservoir states. Then exchange source validity and test whether the obsolete record can be depressed and the new one acquired. Controls: outcome disconnected, trace disconnected, records equalised, and a uniformly stronger negative writer. Apply the physical transplant of test 1 to the learned populations to test whether both source preference and memory influence follow the reservoir state.

Before attempting a multi-link network, rerun the model at the measured device parameters: linewidths, shifts per excitation, reservoir and trace lifetimes, maximal shift, graded update sizes and diffusion length.

S9 Numerical accuracy, data and code

S9.1 Integration step and the result markers

The G2 and G3 results use the reference step $\Delta t = 0.002$; coupled refinements use 0.001 and 0.0005 with matched noise. A verdict is called robust when it is unchanged over the tested range of unresolved networks; this does not mean that every trajectory or count is identical. The step comparison uses short segments of the reference histories at the three steps, an order test of whether differences decrease with the step, and a long deterministic replay; the segment identifiers and tolerances are in the deposit. Table S15 gives the status of each trace rule and the meaning of the markers used in this Supplement.

Table S15: Numerical status of the G2 and G3 results under step refinement.

Rule	Status at $\Delta t = 0.002, 0.001$ and 0.0005
Use ([†])	Some short-segment quantities are not fully resolved at the reference step. The G2 formation and co-read verdicts, the redraw failure, and the G3 durability and reader-withdrawal verdicts are the same at all three steps; G3 durability is 21/24, exactly at the target; the fresh false-record count can change from zero to one and the G3 persistent-false count from one to two, without a change of verdict.
Co-presence ([†])	The reference-step segment target is not met in G2 and some G3 quantities are unresolved, but the aggregate verdicts do not change: the durable-network count is zero at all three steps, the fresh false-record count is 34 at the reference step and 32 under refinement, and the co-read count spans 28–32 rather than 33.
Admission	Meets the reference-step segment target, and its G3 durability passes all reported refinement checks. The fresh transplant count is 21 at the reference step and 20 under refinement.
FML1	Separate re-implementation at $\Delta t = 0.002$; faster acquisition than at $\Delta t = 0.0005$, where its sampled formation-time distribution is consistent with that of the original code, whose timing does not change with the step; equivalence is not established (Section S7.10).
FRCP1	$\Delta t = 0.0005$; no further refinement (Section S7.8).
S3VAL, RLR1, VAL1	Original code at $\Delta t = 0.0005$; halving the step on the six development networks under the three VAL1 writers (PC, F1 and U) left all 18 complete-sequence and 108 criterion labels unchanged, with event times moved by at most one cycle and one dose by 0.73% (RES1, below); a two-step sensitivity check, not a convergence order. The 60-network cohort was run at $\Delta t = 0.0005$ only. Agreement of the two scorers on every label is a check of scoring, not of numerical convergence (Section S7.11).

Step halving on the development networks (RES1). A separate registered study repeated the six development networks N1–N6 under all three VAL1 writers at $\Delta t = 0.0005$ and 0.00025 , with noise increments coupled so that the fine run received the same realised noise as the coarse run; both scorers evaluated every history and each was reconstructed independently. No functional label changed: the complete-sequence label was the same at both steps in 18 of 18 histories and the criterion labels in 108 of 108 comparisons; recorded events moved by at most one cycle, and one dose changed by 0.73%. These are the exposed development networks, so their counts are descriptive and are not pooled with the 60-network cohort. The check is a two-step consistency test with one refinement factor: it establishes neither a convergence order nor continuous-time certification, does not cover the fresh cohort, and does not alter the VAL1 result. The identity checks, segment runs, per-network coarse/fine matrix and archive mapping are in the deposit.

S9.2 Replay and independent implementations

In the G1 audit, the historical runs of the original C implementation of R1–R5 were reproduced byte for byte after the order of the two Gaussian draws used for complex field noise was fixed, and independent right-hand-side and integration-step checks agreed with that code. The separate Python implementation of R3 was rescored from saved logs; its noise generator was not controlled by the task seed, so those histories could not be replayed path by path. It reproduces model choices absent from the printed specification, so it is not an independent reconstruction. The FML1 simulator was written from the published equations; its checks against this Supplement are given in Section S7.10. Article Fig. 2a,b recomputes one of the 24 G2/G3 admission histories over cycles 1–750 with the original simulation code. An added observer records the state within selected cycles without drawing random numbers or changing the state. Despite a different compiler and processor, the recomputed per-cycle log matches all 750 original lines, and end records, sensing signals and selection probabilities agree within 2×10^{-14} . Figure S2 shows three cycles of this history.

Separate algebraic and numerical assertions checked the analytical results of Sections S2–S6; the scripts will be deposited. These checks address the stated analytical problems and are not new network runs.

S9.3 Data and code

Table S16 lists the items of the data and code deposit and the sections they support.

Table S16: Items of the data and code deposit. A repository identifier will be added when the archive is released.

Item	Content	Supports
Network simulations R1–R5, G1–G3	Simulation source, run logs, evaluators, registered criteria and scoring	Sections S7.1–S7.4, S9.1
Functional-control calculations	Fixed-record reader datasets, trace and bound calculations, inequality checkers	Sections S2.3, S3.3–S3.4
FPL1 and FRCP1	Histories, event records, checkpoints, arm outputs, analysis code and independent recomputations	Sections S7.7, S7.8
Article Fig. 2a,b	Observer patch, rerun commands and the recomputed history	Section S9.2
Sensing lifetimes	Initial states, continuation outputs, per-state survival and boundary tables	Section S7.9
FML1	Independent simulator, run scripts, analysis package, measured event classes, per-cycle logs of all arms (including the capacity-alone and neighbouring arms), trajectory figures and the step checks	Sections S6, S7.10
Capacity writer	Scalar implementation and deterministic tests of the capacity factor	Section S6
LMR2 and S3VAL	Registrations, scorers and their agreement, item matrices for all five writer laws, raw histories and assay records	Section S7.11
RLR1	Candidate specification, premise check, reset and late-activation runs, six-network development histories and scores	Section S7.12
VAL1 and RES1	Design lock, pre-written outcome texts, per-history and aggregate records of both scorers, H-labelled networks, class arrays, component tables, bootstrap intervals, secondary readouts, step-halving histories and the independent audit	Sections S7.13, S9.1
Supplement archive	The passages, tables and figures removed from this Supplement during the revisions, as compiled documents with their original study identifiers, including the historical registry table, the prepared-state monitoring campaign, the one-sided writer analysis, the FML1 trajectory figure and full tables, and the VAL1 bookkeeping	All sections

References

- [1] Kalinin, K. P. & Berloff, N. G. Polaritonic network as a paradigm for dynamics of coupled oscillators. *Phys. Rev. B* **100**, 245306 (2019).
- [2] Yagishita, S. *et al.* A critical time window for dopamine actions on the structural plasticity of dendritic spines. *Science* **345**, 1616–1620 (2014).
- [3] Brzosko, Z., Schultz, W. & Paulsen, O. Retroactive modulation of spike timing-dependent plasticity by dopamine. *eLife* **4**, e09685 (2015).
- [4] Falk, M. J., Strupp, A. T., Scellier, B. & Murugan, A. Temporal contrastive learning through implicit non-equilibrium memory. *Nat. Commun.* **16**, 2163 (2025).
- [5] Goldt, S. & Seifert, U. Stochastic thermodynamics of learning. *Phys. Rev. Lett.* **118**, 010601 (2017).
- [6] McCloskey, M. & Cohen, N. J. Catastrophic interference in connectionist networks: the sequential learning problem. *Psychol. Learn. Motiv.* **24**, 109–165 (1989).
- [7] French, R. M. Catastrophic forgetting in connectionist networks. *Trends Cogn. Sci.* **3**, 128–135 (1999).
- [8] Schmidt, D. *et al.* Tracking dark excitons with exciton polaritons in semiconductor microcavities. *Phys. Rev. Lett.* **122**, 047403 (2019).
- [9] Stepanov, P. *et al.* Dispersion relation of the collective excitations in a resonantly driven polariton fluid. *Nat. Commun.* **10**, 3869 (2019).
- [10] Rozas, E. *et al.* Polariton-dark exciton interactions in bistable semiconductor microcavities. *Phys. Rev. B* **108**, 165411 (2023).
- [11] Ohadi, H. *et al.* Spontaneous spin bifurcations and ferromagnetic phase transitions in a spinor exciton-polariton condensate. *Phys. Rev. X* **5**, 031002 (2015).
- [12] Gavrillov, S. S. *et al.* Bistability and nonequilibrium transitions in the optically polarized system of cavity polaritons under nanosecond-long resonant excitation. *Phys. Rev. B* **85**, 075319 (2012).
- [13] Kurdyubov, A. S. *et al.* Optical control of a dark exciton reservoir. *Phys. Rev. B* **104**, 035414 (2021).
- [14] Suchomel, H. *et al.* Prototype of a bistable polariton field-effect transistor switch. *Sci. Rep.* **7**, 5114 (2017).

- [15] Hayat, A. *et al.* Dynamic Stark effect in strongly coupled microcavity exciton polaritons. *Phys. Rev. Lett.* **109**, 033605 (2012).
- [16] Demirağ, Y. *et al.* PCM-trace: scalable synaptic eligibility traces with resistivity drift of phase-change materials. In *2021 IEEE International Symposium on Circuits and Systems (ISCAS)*, 1–5 (2021).
- [17] Wang, Z., Marandi, A., Wen, K., Byer, R. L. & Yamamoto, Y. Coherent Ising machine based on degenerate optical parametric oscillators. *Phys. Rev. A* **88**, 063853 (2013).
- [18] Takata, K., Marandi, A. & Yamamoto, Y. Quantum correlation in degenerate optical parametric oscillators with mutual injections. *Phys. Rev. A* **92**, 043821 (2015).