

Evidence-coupled bistability converts dependency structure into a recovery geometry

Hiroki Saito

hiroki.saito.neuro@gmail.com

Department of Neuropsychiatry, Nippon Medical School <https://orcid.org/0000-0002-4964-8559>

Research Article

Keywords:

Posted Date: September 11th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10620384/v2>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

**Evidence-coupled bistability converts dependency structure into a recovery
geometry**

Hiroki Saito^{1,2}, M.D., Ph.D.

¹Onda-Daini Hospital, Matsudo-shi, Chiba, Japan

²Department of Neuropsychiatry, Graduate School of Medicine, Nippon Medical
School, Sendagi 1-1-5, Bunkyo-ku, Tokyo 113-8602, Japan

* Corresponding author: Hiroki Saito

Department of Neuropsychiatry, Graduate School of Medicine, Nippon Medical School,
Sendagi 1-1-5, Bunkyo-ku, Tokyo 113-8602, Japan

E-mail: hiroki.saito.neuro@gmail.com

PHONE: +81-3-3822-2131 FAX: +81-3-5814-6287

Abstract

Dependency structure is known to constrain the order in which an adaptive system regains its capacities, but not the rest: whether recovery is self-sustaining, is held without support, stalls, or collapses. Here we place an evidence-coupled bistable cascade on a finite forward acyclic dependency chain: each capacity carries an evidence-dependent threshold, evidence accrues only while every prerequisite is online, and accumulated evidence lowers that threshold. We prove that consolidated evidence accelerates every capacity super-linearly, that the recovered and failed configurations are two stable states, and that the surface on which the root switches regime is exactly invariant, its position independent of the evidence timescale. That invariance is not bought by a postulated selector: the update follows from a two-regime Bayesian model with ambiguity-suspended updating, exclusivity arising in the sharp limit. Directional coupling then fixes the outcome: above an explicit sufficient bound every trajectory off the surface converges to full recovery or full collapse, while sufficiently weak coupling admits a stable partial state at consolidated evidence. Dependency structure thus gates when evidence can accumulate, and that gating, coupled to bistability, yields a recovery geometry whose outcomes depend on coupling strength and initial condition. Catatonia is the first biological instantiation.

Introduction

Systems built from interdependent capacities lose and regain them in constrained ways.

A nervous system after injury, an immune compartment after depletion, a learner acquiring a curriculum and a patient recovering from a psychomotor syndrome share a feature: some capacities cannot be exercised, or defined, until others are in place. That structure is known to constrain the order in which capacities return; what it constrains beyond order is not. Whether recovery becomes self-sustaining, whether the recovered configuration is held without external support, whether a system stalls in a partial configuration time does not resolve, and whether an improving system can convert abruptly into whole-system failure are questions about outcomes, which the sequence does not answer.

Two literatures answer half the question each. Dependency-order theory takes a graph of prerequisite relations¹ and asks which capacity may precede which; under gating axioms the admissible orders are its linear extensions, whatever the rates²⁻⁴. Such results constrain sequence, not outcome. Bistable dynamics answers the complement: a capacity with an Allee-type threshold has a failed and a functional state, and which one a trajectory approaches is fixed by the basin it starts in⁵⁻⁹. Such results are indifferent to prerequisite structure, one component's threshold not usually depending on whether the

components it needs are online.

Recovery-order theories answer which capacity may precede which, bistable models which basin a system approaches. Neither explains how prerequisite structure reshapes basin geometry through accumulated evidence. That is the gap addressed here, and the coupling is not symmetric. Dependency structure gates when evidence can accumulate, the evidence rate being the joint availability $C(r) = \prod_j r_j$, which vanishes while any prerequisite is offline; accumulated evidence then deforms the bistable threshold shared by every capacity. Coupled to that bistability, the gating generates a recovery geometry whose outcomes depend on coupling strength and initial condition, so the structure constrains not only the order of recovery but whether a system consolidates, stalls or collapses.

The question is therefore: once a dependency structure is known, what dynamical architecture explains self-sustaining recovery, consolidation, relapse, stable partial recovery and catastrophic collapse within one system? The answer proved here is that an evidence-coupled bistable cascade on the chain suffices, strong coupling yielding a global dichotomy, whereas sufficiently weak coupling admits a persistent partial state.

Catatonia made the problem concrete and remains the instantiation carried through the

paper. Recovery there can be observed directly and, in most patients, rapidly^{10,11}, and is documented in enough detail that its shape, not only its endpoint, is a matter of published record^{11,12}. Four dynamical features have been reported separately and none has explained them together: improvement is a staged reorganisation after which further escalation is unnecessary^{12,13}; the recovered state is then maintained at no continuing control cost, benzodiazepines often being tapered without relapse, though tapering returns the full state in a subset of patients rather than producing graded worsening^{14,15}; and an improving state can convert abruptly into a malignant one, a whole-system collapse rather than an intensification¹⁶⁻¹⁹. A fifth observation constrains any account of failure: not every unrecovered patient collapses, some stabilising in a chronic state with basic reactivity present but organised action, spontaneous initiation and contextual coherence absent^{13,20}.

These reports are independent of one another and of this work. Together they say that recovery is not a monotone climb along a severity axis but a passage between distinct configurations, with a threshold crossable in either direction and a stable outcome that is neither recovery nor failure. The order in which capacities return is not re-derived here: in companion work that sequence, the Saito Loop, is obtained from the zero-pattern of the joint Fisher information as the unique prerequisite order over sensory, policy,

motivational, fast-volatility and slow-volatility precision, and a graph-general theory fixes the admissible orders for any dependency graph²⁻⁴. Those papers leave the geometry open; this one determines it.

One caveat is stated at the outset. Identifying the slow coordinate of a system with accumulated Bayesian evidence is in general a modelling commitment, not a mathematical consequence; it needs an observation model defining that coordinate as a monotone function of the posterior precision of a slow environmental model²¹. For catatonia that model is available and the identification is derived from it below; elsewhere it must be supplied. The value of the construction is not that it is forced by observation, but that a small set of structural assumptions generates the outcome repertoire without introducing a backward edge into the dependency structure, and that removing each listed structure eliminates at least the corresponding component (Table 1).

Results

The object proved here is one structure, an evidence-coupled bistable cascade on a forward acyclic dependency chain (Fig. 1); T1 to T5 are its properties, not five separate

facts. They are stated for the displayed chain of four capacities with one evidence variable and proved in the Methods, by induction over the chain, for a chain of any finite length. Branching structures, in which a capacity has several parents, would need the coupling summed over parents and switching surfaces at several roots; they are not proved here.

An evidence-coupled cascade on a dependency chain

Capacities r_k in $[0,1]$, $k = 0$ to 3 , are indexed along a forward acyclic dependency chain, r_0 the root, and the evidence variable Q in $[0,1]$ is dimensionless. Writing $z = r_0 - a(Q)$, the dynamics are

$$dr_0/dt = g(r_0, Q) [+ u(t), \text{ an external excursion applied to the root }]$$

$$dr_k/dt = g(r_k, Q) + \kappa_k (r_{k-1} - r_k), k = 1 \text{ to } 3$$

$$dQ/dt = \varepsilon [\chi_+(z) C(r) (1 - Q) - \rho \chi_-(z) Q]$$

$$g(r,Q) = r (1 - r) (r - a(Q)), C(r) = \prod_{j=0}^3 r_j, a(Q) = a_0 - (a_0 - a_1) Q^p, p > 1.$$

The evidence coordinate is $Q = 1 - \exp[-(\Lambda - \Lambda_0)/\Lambda^*]$, a monotone coordinate on the accumulated Fisher information Λ of the slow model. $Q = 1$ corresponds to $\Lambda \rightarrow \infty$, so the recovered state is a compactified consolidated limit, not a claim that infinite

precision is reached in finite time, and Q stays below 1 at all finite times. The construction is fixed by four structural axioms, four generative premises and one observation model.

Premise (observation model). The evidence variable is a monotone coordinate on the posterior precision of a slow environmental model, the analogue of commitments M1 to M3 of the companion theory. In the catatonia instantiation the premise is not assumed but derived from that theory's observation model, as the closure result below shows.

A (dependency structure). The capacity layer is a forward acyclic diffusive chain from r_{k-1} to r_k with non-negative coupling, hence a cooperative, order-preserving cascade^{5,6}, and the root has no incoming coupling ($\kappa_0 = 0$). No edge runs backwards, so the chain carries the prerequisite order supplied by the companion theory.

B (bistability). Each capacity carries an evidence-dependent Allee threshold^{7,8}, $g(r, Q) = r(1-r)(r-a(Q))$, giving a failed and a functional basin.

C (evidence generation). Consistency of the chain generates valid Fisher information²² for the slow model only when every prerequisite is online: $d\Lambda/dt = I_0 C(r)$, $C(r) = \prod_j r_j$.

D (basin deepening). Accumulated evidence lowers the shared threshold, decreasingly and concavely: $a'(Q) < 0$ and $a''(Q) < 0$. Concavity is what makes the self-catalysis

super-linear. Neither sign is posited: Axiom D is derived conditionally from a specified active-inference decision layer²³⁻²⁷. The negative slope follows in its informative and non-degenerate regime; concavity holds exactly when an explicit inequality relates the precision sensitivity to the scale of the evidence coordinate (Supplementary Methods).

Regime-selective updating is a limit, not an axiom

E1 (competing regimes). The slow model is at all times in one of two mutually exclusive regimes, $m \in \{M_{\text{cont}}, M_{\text{reset}}\}$: the environmental model is retained and its evidence accumulated, or rejected and its evidence reset, which is the change-point reading²⁸.

E2 (evidence margin). The root margin $z = r_0 - a(Q)$ is, up to a positive constant γ , the log posterior odds between the regimes: $\log[P(M_{\text{cont}} | y_{0:t}) / P(M_{\text{reset}} | y_{0:t})] = \gamma z$, a modelling commitment of the same kind as the observation-model premise.

E3 (responsibility-weighted update). Each regime carries its own precision law, $dQ/dt = \varepsilon C(r)(1 - Q)$ under M_{cont} and $dQ/dt = -\varepsilon pQ$ under M_{reset} , combined by Bayesian model averaging with the responsibilities $w_+(z) = 1/(1 + e^{-\gamma z})$ and $w_- = 1 - w_+$, so that $\log[w_+/w_-] = \gamma z$ exactly as E2 requires²⁹.

E4 (ambiguity-suspended accumulation). Evidence is credited to a model only so far

as that model is identified: the information rate carries the factor $\lambda_\gamma(z) = L(w_+(z))$, with L continuous on $[0,1]$, $L(1/2) = 0$, $L(w) > 0$ for $w \neq 1/2$ and $L(0) = L(1) = 1$. The member used below is the normalised entropy deficit $L(w) = 1 - H(w)/\log 2$, H the binary entropy in nats; L^2 and $|2w - 1|$ are two others. Writing λ as a function of the posterior rather than of z makes the class uniform as $\gamma \rightarrow \infty$. E4 is an added information-theoretic principle, not a consequence of Bayes rule, which at equal odds averages rather than suspends the update. Suppressing updating while the model is unidentified is the ambiguity-dependent learning rate of Bayesian volatility models^{26,27,30,31}, and a violation of size η at the boundary moves the switching surface by $O(\epsilon\eta)$ (Supplementary Methods).

Proposition 1 (finite sharpness and its limit). E1 to E4 give $dQ/dt = \epsilon[\chi_+(z) C(r)(1 - Q) - \rho \chi_-(z) Q]$ with $\chi_\pm(z) = w_\pm(z) \lambda(z)$, and two parts separate. (a) For every finite γ and every L in the class of E4, $\chi_+(0) = \chi_-(0) = 0$, which by Proposition 2 is exactly what makes $S = \{r_0 = a(Q)\}$ invariant: the switching surface does not require the sharp limit. (b) As $\gamma \rightarrow \infty$, $\chi_+ \rightarrow 1_{[z>0]}$ and $\chi_- \rightarrow 1_{[z<0]}$, uniformly on $\{|z| \geq \delta\}$ for every $\delta > 0$ and every L in the class, since $w_+ \rightarrow 0$ or 1 uniformly there and L is uniformly continuous on $[0,1]$; the supports become disjoint, and it is exclusivity, not invariance, that the closed-form evidence law of T4 and the monotonicity of z used in T5 require. If E4 is

dropped and the regimes are averaged, then $\chi_{\pm}(0) = 1/2$, S is not invariant and the boundary drifts with the evidence timescale. Proof in Methods.

E (continuous regularisation of the exclusive limit). The limit pair of Proposition 1(b) is discontinuous at $z = 0$ and is therefore not used as the vector field. It is replaced by a one-sided regularisation with the same structure: $\chi_+ \geq 0$, zero for $z \leq 0$ and positive for $z > 0$, symmetrically χ_- , both locally Lipschitz, continuously differentiable on the interiors of their active half-lines, and bounded below by a positive constant on compact subsets of those half-lines. Continuity alone would not do: it gives existence but not uniqueness, and the Jacobians of T2 need differentiability at the corners, which lie away from the switching surface. This is a choice, not a deduction; no functional form enters, T1 to T5 using exactly the displayed properties (Methods).

Scope of the general claim. What is proved holds for finite forward acyclic dependency chains with non-negative coupling, evidence generated by joint availability of every capacity, evidence-dependent bistable thresholds satisfying Axioms B and D, selectors in the class E and, where stated, the coupling condition of T5. What is not claimed is listed in the Supplementary Methods; Table 2 records the status of every result.

Evidence accumulation and the acceleration of recovery

T4. Let $z(0) > 0$, or restrict to a time interval on which $z(t) > 0$, and let $N(t) = \int_0^t \chi_+(z(s)) C(r(s)) ds$ be the accumulated valid stable evidence. Then $Q(t) = 1 - (1 - Q_0) \exp[-\varepsilon N(t)]$. For $Q_0 < 1$, Q is strictly increasing in N and $Q(t) < 1$ at every finite time; for $Q_0 = 1$ it is constant (Supplementary Methods). Because $\{z > 0\}$ is forward invariant (T3(c1)) the interval is not bounded: a trajectory with $z(0) > 0$ stays in the recovery regime. On $\{z < 0\}$ the counterpart is pure attenuation, $Q(t) = Q_0 \exp[-\varepsilon \rho \int_0^t \chi_-(z(s)) ds]$. The closed form is not asserted where both channels act, as at finite γ ; the general linear solution is in Methods. That Q is the posterior precision of a slow environmental model is the observation-model premise; what is proved is the law of its time evolution.

T1. Consolidated evidence raises the recovery rate of every capacity, $\partial(dr_k/dt)/\partial Q = -r_k(1-r_k) a'(Q) > 0$; since downstream recovery raises the evidence rate $C(r)$, this is evidence-mediated downstream-to-upstream acceleration. The response is convex in Q , $\partial^2(dr_k/dt)/\partial Q^2 = -r_k(1-r_k) a''(Q) > 0$, and simultaneous downstream recovery is super-additive, $\partial^2 C/\partial r_i \partial r_j = \prod_{l \neq i,j} r_l > 0$ for $i \neq j$. All three follow by direct differentiation of g and C , using $a' < 0$, $a'' < 0$ and $0 < r_k < 1$ (Supplementary Methods); what is proved is convexity of the map from Q to dr_k/dt , not of the trajectory in time.

Two stable configurations and an exact switching surface

T2. The recovered state $E_{\text{rec}} = (1,1,1,1,1)$ is an exact equilibrium, hyperbolically asymptotically stable for every $\kappa_k \geq 0$, $\varepsilon > 0$ and $\rho \geq 0$. For $\rho > 0$ the failed state $E_{\text{path}} = (0,0,0,0,0)$ is an exact, isolated, hyperbolically asymptotically stable equilibrium. At $\rho = 0$ the segment $\{(0,0,0,0,Q)\}$ is a continuum of equilibria, so E_{path} is not isolated: Lyapunov stable in the capacity coordinates, not asymptotically stable in the full system. Proof in Methods.

T3(a) exact frozen-evidence trigger. With Q fixed, an excursion leaving $r_0 = 1 - \Delta$ sends the root below threshold exactly when $\Delta > \Delta_c(Q) = 1 - a(Q)$, with $Q_c = [(a_0 - (1-\Delta))/(a_0 - a_1)]^{1/p}$. The root field has the single interior unstable zero $r_0 = a(Q)$ and no incoming coupling, which makes the criterion exact. The tolerated excursion grows as evidence consolidates, which is the sense in which the fragile window closes.

T3(b) propagation (sufficient coupling). If the root falls and $\kappa_k(a_1 - \eta) > M$ for some $0 < \eta < a_1$ (M the Allee hump height), the whole chain collapses in the forward order, and Q to 0. The bound is sufficient, not necessary: some coupling is needed, since with $\kappa_k = 0$ a fallen root leaves the downstream capacities in place, but propagation occurs well below it. Proof in Methods.

T3(c1) exact invariant switching surface. The surface $S = \{r_0 = a(Q)\}$ is invariant for all ε and ρ , and the regions $z > 0$ and $z < 0$ are forward invariant, so $\Delta > \Delta_c(Q)$ holds if

and only if the root enters and remains in the collapse regime. S is the exact, timescale-independent switching surface of the root, and is not asserted to be the global basin boundary. Proof in Methods.

Proposition 2 (necessity of boundary-zero updating for exact switching). Suppose the evidence update has the separable form $dQ/dt = \varepsilon[\alpha(z)C(r)(1 - Q) - \rho\beta(z)Q]$, with α and β continuous and $\alpha'(Q) \neq 0$, and that S is invariant for every $\varepsilon > 0$ and $\rho > 0$. Then $\alpha(0) = \beta(0) = 0$, and conversely that condition makes S invariant. Boundary-zero updating is therefore what the timescale-independent manifold of T3(c1) requires, and Proposition 1(a) supplies it at every finite sharpness. Exclusivity is strictly stronger: not needed for invariance, but used by the one-sided law of T4 and the monotonicity argument of T5, and supplied in the limit by Proposition 1(b). Proof in Methods.

T3(c2) whole-chain outcome from a recovered configuration. Fix $r_k(0+) = 1$ for $k = 1, 2, 3$ and $r_0(0+) = 1 - \Delta$. Then $\Delta > \Delta_c(Q)$ with the bound of T3(b) gives whole-chain collapse, and $\Delta < \Delta_c(Q)$ absorption back to E_{rec} , unconditionally in ε and ρ . Absorption is not unconditional over arbitrary data: a root on the recovery side with downstream capacities failed need not recover them at low coupling. Proof in Methods.

Removing each listed structure eliminates at least the corresponding component (Table

1, Fig. 2); some removals affect more than one, so the decomposition is minimal only in the operational sense used here.

Strong coupling yields a global recovery-collapse dichotomy

The results so far do not require global bistability. Whether every trajectory off the switching surface converges to a corner needs an extra condition on directional coupling.

T5 (global two-outcome convergence under strong coupling). Assume $\rho > 0$, selectors in the class of E, and $\kappa_k > K^* := \max \text{ over } a \in [a_1, a_0] \text{ of } (1 - a + a^2)/3$, $k = 1, 2, 3$. Then $z(0) > 0$ implies $(r, Q) \rightarrow E_{\text{rec}}$ and $z(0) < 0$ implies $(r, Q) \rightarrow E_{\text{path}}$, for every initial condition in the state cube, the endpoints $r_0(0) = 1$ and 0 included. The set $z = 0$ is invariant and contains the non-attracting equilibrium continuum $E_S = \{(a(Q), a(Q), a(Q), a(Q), Q)\}$, to which every orbit within S converges at its frozen evidence level. The two regions are therefore exactly the basins, S their common boundary, and no other compact invariant set attracts a trajectory off S. K^* is sufficient, not sharp. For all sufficiently small positive equal coupling a locally asymptotically stable mixed branch exists, the root recovered and downstream capacities low, so the coupling condition cannot be dropped; there $\kappa_1 < a_1^2/4$ is the exact condition for two distinct low first-stage roots, the lower hyperbolically stable, coalescing in a saddle-

node at $\kappa_1 = a_1^2/4$. Proof in Methods.

T5 is exact in the regularised exclusive limit E. At finite sharpness S is still invariant, so the capacity-level dichotomy survives exactly while the recovered evidence coordinate is displaced by at most $\rho \exp[-\gamma(1 - a_0)]$ (Supplementary Methods).

With the illustrative parameters ($a_0 = 0.60$, $a_1 = 0.15$) the threshold is $K^* \approx 0.291$ and the illustrative coupling $\kappa = 0.6$ satisfies it; the first-stage low-root condition is $\kappa_1 < a_1^2/4 \approx 0.0056$. The term global attractor is avoided, the invariant set including the equilibrium continuum on S.

The two directions close on one another (Fig. 2). Forwards, consistency of the chain (Axiom A) generates evidence (Axiom C, T4), deepening the basins (Axiom D), which accelerates recovery and stabilises the recovered state (T1, T2). Backwards, a threshold crossing at the root enters the change-point regime (E), attenuating evidence and, under sufficient coupling, collapsing every capacity (T3).

Weak coupling permits stable partial recovery

Sufficiently weak coupling gives a third outcome, specific enough to be looked for.

Prediction 1. Weak directional coupling produces not mild residual deficits but a severe stationary partial state. At the illustrative parameters with $\kappa = 0.003$ the stable mixed

equilibrium is $r \approx (1.00, 0.024, 0.0005, 0.00001)$ at $Q = 1$: root restored, downstream capacities absent, evidence consolidated. Its state-block eigenvalues are of order 10^{-1} while the evidence direction is effectively neutral, so it is entered on the ordinary recovery timescale and persists. The signature is a profile, not a severity score.

Prediction 2. Coupling above the bound produces all-or-none recovery, with no stable intermediate.

Prediction 3. Because the mixed state persists at fully consolidated evidence ($Q = 1$), raising evidence without restoring directional coupling cannot remove it; only $\kappa > K^*$ eliminates the branch, so time and further consolidation alone are predicted not to resolve it.

Catatonia as a biological instantiation

In the catatonia instantiation the capacities are sensory, policy, motivational and fast-volatility precision, the evidence variable is the accumulated identifiability of slow-volatility precision, and the chain is $\pi_s \rightarrow \beta \rightarrow \pi_m \rightarrow \pi_{v_fast} \rightarrow \pi_{v_slow}$, obtained in the companion computational study from the joint Fisher information³. Under this identification the four reported clinical features correspond respectively to four proved properties of the model: self-sustaining improvement to T1, maintainable remission to

T2, withdrawal-precipitated relapse to T3 read at low consolidated evidence, and drug-precipitated malignant conversion to T3(b). The fifth observation, chronic partial catatonia, is the weak-coupling branch of T5, whose predicted profile is restored reactivity with organised action, spontaneous initiation and contextual coherence absent²⁰.

The present model does not inherit the companion recovery-order axioms, which require a gate vanishing whenever a prerequisite is offline, a deficit-closing recovery function, and a common impaired onset⁴. It satisfies none: the coupling is diffusive, so at $r_{k-1} = 0$ the drive of capacity k is $g(r_k, Q) - \kappa_k r_k$, which does not vanish; the field is negative below threshold, so not deficit-closing; and under a uniform onset the diagonal is invariant, so the capacities recover synchronously (Table 3). That theorem therefore does not apply here, the order being reproduced instead by forward diffusive coupling, which prevents a descendant from overtaking its prerequisite, with the root-specific input $u(t)$ supplying the symmetry breaking a uniform onset does not.

Proposition 3 (order preservation under root-initiated recovery). With $u(t) = 0$ after initiation and $r_0(0) \geq r_1(0) \geq r_2(0) \geq r_3(0)$, the order is preserved for all $t \geq 0$, so at any common threshold q the first-passage times satisfy $T_0(q) \leq T_1(q) \leq T_2(q) \leq T_3(q)$ (Supplementary Methods). It covers the four capacity variables only; Q is ordered

against them only under the closure reading below.

What the construction adds is not the order but the geometry, and in this instantiation the observation-model premise is a consequence rather than an assumption.

Proposition 4 (closure of the loop). In the companion generative model³ the slow-volatility channel is the only one whose likelihood depends on $\pi_{V_{\text{slow}}}$, and its expected count is proportional to the joint availability of the four capacities below it. Its statistic is the set of per-context estimates of $\pi_{V_{\text{fast}}}$, usable only when every capacity is online, so the information accumulated by time t is $I_0 \int_0^t C(r) ds$ and $Q = 1 - \exp[-(\Lambda - \Lambda_0)/\Lambda^*]$ is a strictly increasing bijection of it onto $[0,1)$. The first-level channel has the same zero set and $C(r)$ dependence, but a prefactor varying with $\pi_{V_{\text{slow}}}$, so it corroborates the gate rather than supplying the law. The evidence variable of this paper and the statistic that identifies the top node of the companion chain are therefore the same object.

The loop is closed and not circular, within the stated observation and decision model classes. Joint availability is the only source of information about $\pi_{V_{\text{slow}}}$, that information is Q , and Q lowers the shared threshold, which raises availability, the return leg being the decision rule behind Axiom D. No cycle is added to the capacity layer, Q not being a node of the chain. For slow enough evidence the fifth milestone is reached last, the

reported clinical order (Supplementary Methods). Table 3 gives the division of labour across the four studies.

Discussion

The principle established here is that dependency structure constrains more than sequence, and the route is specific. It gates when evidence can accumulate: the evidence rate is the joint availability $C(r) = \prod_j r_j$, so a single offline prerequisite halts accumulation however well the rest of the system is doing. Accumulated evidence then deforms the threshold shared by every capacity, moving the stable configurations and the surface between them. Coupled to evidence-dependent bistability, that gating generates a recovery geometry in which coupling strength and initial condition fix whether a system consolidates, stalls or collapses. The word geometry is deliberate and landscape avoided: no potential function is claimed, the system not being a gradient flow, so what is proved concerns invariant surfaces, basins and their boundary rather than the depth of a well.

What is proved is summarised in Table 2: a convex evidence-to-rate gain (T1); two hyperbolic equilibria, the failed one needing an active attenuation channel (T2); an exactly invariant switching surface, hence a threshold independent of the evidence

timescale (T3(c1)), which Proposition 2 shows to need boundary-zero updating and nothing stronger; a tolerated excursion widening as evidence accumulates (T3(a)); and, above an explicit sufficient bound, two basins with S as their boundary, whereas sufficiently weak coupling admits a stable partial configuration (T5). K^* is sufficient, not sharp, and the intermediate range is unclassified.

The relation to recovery-order theory is a division of labour, not an overlap. A clinical study reports the order in a cohort²; a computational study derives the dependency structure from an active-inference architecture, its unique linear extension being that order³; a graph-general theory proves that for any graph under gating axioms the admissible orders are exactly its linear extensions⁴. All three stop at order. The present paper takes the structure as premise and asks what dynamics produce the observed outcomes; the order it reproduces (Proposition 3) is a by-product of forward coupling and root-initiated onset, in a different axiom class (Table 3). That theory says which orders a structure permits, this one which endpoints, once thresholds are evidence-dependent.

Two constructions are disclosed. The saturating coordinate $Q = 1 - \exp[-(\Lambda - \Lambda_0)/\Lambda^*]$ makes the recovered state an exact compactified equilibrium, needing no parameter condition; the second, the regime decision, is load-bearing. Its three steps differ: the

finite-sharpness selector of E1 to E4 already makes S invariant, the identifiability weighting of E4 forcing $\chi_{\pm}(0) = 0$; the sharp limit adds exclusivity, and with it the closed-form law of T4 and the monotonicity of z used by T5; and since that limit is discontinuous, a one-sided regularisation keeps the problem well posed. Only the first step is inference: E4 is an added principle, not a consequence of Bayes rule, and the last a choice. The residual arbitrariness is small: no proof uses the member of the E4 class or the shape of the regularisation, and a violation of the boundary zero by η moves the surface only by $O(\epsilon\eta)$. Dropping the weighting altogether costs invariance: $\chi_{\pm}(0) = 1/2$ and the boundary drifts with the evidence timescale, from 0.5940 at $\epsilon = 0.02$ to 0.5786 at $\epsilon = 0.00125$, approaching 0.5774 as the evidence is made slow (Table 1). Standard Bayesian change-point inference is soft, carrying a posterior over run lengths²⁸, so the finite- γ model is the honest object and the sharp system its boundary case, differing by an exponentially small displacement of the recovered corner.

Stable partial recovery is a regime of the construction, not a failure: an incomplete outcome is not a slow trajectory towards full recovery, nor noise around it, but a hyperbolically stable equilibrium of a weakly coupled cascade, entered on the ordinary timescale. The distinction is testable: a delay resolves with time, a mixed state does not. For catatonia four separately documented features become properties of one cascade,

and a fifth, chronic partial catatonia, marks the weakly coupled regime, not a residue.

Clinical truth is not claimed, nor that the two corners are the only attractors without the strong-coupling condition; the coupling threshold and consolidation time are empirical quantities the theory constrains but cannot supply.

Three other systems have the ingredients; the companion theory asks of two which order is admissible⁴, the questions here concerning destinations and cliffs. The recovery plateau after injury is the first³²: where a proximal capacity precedes distal control and timing, sufficiently weak cross-stage coupling predicts that a plateau is a stable equilibrium with the proximal base restored and the dependent capacities near zero, not a slow approach to full recovery, and further time at unchanged coupling cannot resolve it; above the bound it predicts bimodal endpoints instead of a continuum. Graft failure and late rejection in immune reconstitution are the second³³: since the tolerated excursion $\Delta_c(Q) = 1 - a(Q)$ widens as evidence consolidates, a fixed insult should carry a higher collapse probability early than late, and above the bound should fail whole-compartment rather than gradedly. Catastrophic forgetting in sequential and curriculum learning is the third^{34,35}: interference displacing a foundational skill by more than the tolerated excursion should remove the dependent skills together rather than degrade them smoothly, and be less damaging once the foundation is consolidated. None is

tested here.

The forward problem should not be confused with its inverse. What is proved runs from structure to dynamics: given a dependency chain and evidence-coupled bistable dynamics on it, the outcomes follow. The converse, that observed trajectories identify the structure, is neither proved nor implied; the companion theory makes the same separation⁴. The inverse problem is nevertheless well posed as a programme, and the results identify what such an inference would exploit: timescale independence of the switching surface, widening of the tolerated excursion with accumulated evidence, and an endpoint distribution bimodal above the coupling threshold with a structured partial mode below. These differ from the variance and autocorrelation signatures used to anticipate critical transitions³⁶, which report proximity to a bifurcation, not the structure generating it.

The limitations are the conditions themselves: the theorems hold in the stated class; T5 is exact in the regularised exclusive limit and survives at finite sharpness with the recovered corner exponentially displaced; and it is proved above an explicit sufficient bound while the mixed branch is proved for weak coupling, the intermediate range being unclassified. Branching structures are the natural extension, not proved here: the coupling must be summed over parents, the switching surface defined at each root, and

the triangular induction of T5 redone in a topological order. The construction is falsifiable: a switching boundary moving with the rate of evidence accumulation would count against E4, ambiguous evidence being shared rather than withheld; a partial state resolving with time alone would count against the weak-coupling reading; and a unimodal intermediate endpoint distribution above the bound would count against T5.

In summary, a dependency chain with evidence-coupled bistable dynamics has a determinate recovery geometry: downstream recovery accelerates upstream capacities through accumulated evidence (T1), the recovered and failed configurations are two stable basins (T2), the surface between them is invariant and widens the tolerated excursion as evidence consolidates (T3), and directional coupling separates a dichotomy above the bound from a stable partial state at weak coupling (T5). Catatonia is the first biological instantiation. No new data are analysed here, and the predictions above are the confrontation the theory now needs. There the evidence variable is not a further assumption: the companion observation model makes it the accumulated identifiability of the terminal node.

Methods

Chain length

Everything below is proved for a forward acyclic chain of $n + 1$ capacities, $n \geq 1$: $dr_k/dt = g(r_k, Q) + \kappa_k(r_{k-1} - r_k)$ for $k = 1$ to n with $\kappa_k \geq 0$ and $C(r) = \prod_{j=0}^n r_j$, the root and evidence equations unchanged. The Results, the figures and all numerical work display $n = 3$. Each proof is an induction over k with the root as base and uses only the forward ordering, the non-negative couplings and $0 \leq r_k \leq 1$, so T1 to T5 and Propositions 1 to 3 hold for every finite n ; the one length dependent quantity is the persistence bound $C(r) \geq a_1^{n+1}$, positive for every finite n . Branching structures lie outside this statement.

Model and integration

Integration used SciPy `solve_ivp` (RK45, `rtol` 1e-9, `atol` 1e-11, maximum step 0.5), with solver success checked on every call. The regularisation was $\chi_{\pm}(z) = \exp(-s/z^2)$, zero on the opposite half-line, with $s = 1$ and $s = 0.01$ where a shorter integration was needed; both lie in the class of E and leave every reported outcome unchanged. The finite-sharpness realisation was $\chi_{\pm} = w_{\pm}\lambda$ of Proposition 1 with $\gamma = 20$ and λ the normalised entropy deficit, the members λ^2 and $|2w_+ - 1|$ testing the class-level claim of E4; the plain-averaging ablation drops λ . The 1e-9 cutoff around $z = 0$ is an implementation detail, not part of the selector. Parameters were illustrative ($\kappa = 0.6$, $\varepsilon = 0.02$, $p = 1$, $a_0 = 0.60$, $a_1 = 0.15$, $p = 2$); no search was performed to force any outcome. Quoted values depend on the regularisation and parameter choice; because $\exp(-1/z^2)$ is very small for

small z , timescales are not real times. The code separates results by status and offers three run modes.

Proof of Proposition 1

Under E1 to E4 the posterior responsibilities are $w_{\pm}(z) = 1/(1 + e^{\mp\gamma z}) = (1 \pm \tanh(\gamma z/2))/2$, so $\log[w_+/w_-] = \gamma z$ as E2 requires, and the identifiability weight is $\lambda_{\gamma}(z) = L(w_+(z))$ with L in the class of E4. Averaging the laws of E3 with these responsibilities and weighting the information rate by λ_{γ} gives $dQ/dt = \varepsilon[\chi_+(z)C(r)(1 - Q) - \rho\chi_-(z)Q]$ with $\chi_{\pm} = w_{\pm}\lambda_{\gamma}$.

(a) Finite sharpness. $w_{\pm}(0) = 1/2$ gives $\lambda_{\gamma}(0) = L(1/2) = 0$, hence $\chi_+(0) = \chi_-(0) = 0$ for every finite γ and every L in the class of E4; by Proposition 2 this is precisely the condition under which S is invariant, so the boundary behaviour is delivered at finite sharpness. Only $L(1/2) = 0$ is used, and the three members $L = 1 - H/\log 2$, L^2 and $|2w - 1|$ leave the collapse boundary at Q_c numerically. For finite γ both $\chi_{\pm}$ are strictly positive off $z = 0$, so the supports overlap and the update is a genuine model average: both channels act on every trajectory, the exponential closed form of T4 fails and the general solution of the Supplementary Methods is needed, and z is no longer monotone on $\{z > 0\}$, which the convergence proofs of T5 use. If E4 is dropped, so that the regimes are averaged without the identifiability weighting, then $\chi_{\pm}(0) = 1/2$, the hypothesis of

Proposition 2 fails and the switching boundary acquires an ε dependence, the ablation of Table 1.

(b) Sharp-selection limit. For $z \geq \delta > 0$ one has $w_-(z) \leq e^{-\gamma\delta}$, so $w_+ \rightarrow 1$ uniformly on $\{z \geq \delta\}$; L is continuous on the compact interval $[0,1]$, hence uniformly continuous, so $\lambda_\gamma = L(w_+) \rightarrow L(1) = 1$ uniformly there, and $\chi_+ = w_+\lambda_\gamma \rightarrow 1$ while $\chi_- = w_-\lambda_\gamma \rightarrow 0$. The argument is symmetric on $\{z \leq -\delta\}$ and uses only the class conditions on L , not the particular member, so part (b) holds for the whole class; the convergence is uniform on $\{|z| \geq \delta\}$ for every $\delta > 0$ and the limit pair is $1_{[z>0]}$, $1_{[z<0]}$.

The limit is not the field used. The limit pair is discontinuous at the switching surface: with $\chi_+(0) = 0$ from part (a) the right-hand limit at $z = 0$ is 1, so the pair is continuous on neither closed half-line and would need a solution concept for discontinuous fields.

Premise E replaces it by a one-sided regularisation with the same support, sign, positivity and boundary zero, locally Lipschitz on the line and continuously differentiable on the interiors of the active half-lines. Continuity alone would not suffice: it gives existence but not uniqueness, as $dx/dt = |x|^{1/2}$ shows, and without uniqueness the invariance of S would not keep a trajectory off it; differentiability away from the boundary gives the Jacobians of T2, whose corners sit at $z = 1 - a_1 > 0$ and $z = -a_0 < 0$. The reference realisation $\exp(-s/z^2)$ is smooth and s enters no theorem. This

step is a choice, not a deduction; the theorems use only the stated properties, so no member of the class is privileged. What E1 to E4 do not supply is the number of regimes, the identification of z with the log posterior odds, the identifiability weight, or the sharpness.

Proofs of T2 and T3

T2. At E_{rec} : $g(1, Q) = 0$, $C = 1$ and $z = 1 - a_1 > 0$, so $\chi_- = 0$ and $dQ = \varepsilon\chi_+(1-Q) = 0$ at $Q = 1$; the state block is lower-triangular with diagonal $-(1-a_1) - \kappa_k < 0$, $g_r = 0$ at $r = 1$ decouples Q , and the Q -eigenvalue is $-\varepsilon\chi_+(1-a_1) < 0$ for every ρ . At $r = 0$, $C = 0$ and $z = -a_0 < 0$, so $\chi_+ = 0$ and $dQ = -\varepsilon\rho\chi_-(-a_0)Q$. For $\rho > 0$ this vanishes only at $Q = 0$, an isolated equilibrium with state diagonal $-a_0 - \kappa_k < 0$ and Q -eigenvalue $-\varepsilon\rho\chi_-(-a_0) < 0$. For $\rho = 0$, $dQ \equiv 0$ on $r = 0$, so every $(0, 0, 0, 0, Q)$ is an equilibrium and the corner is non-isolated. E_{rec} is therefore hyperbolic irrespective of ρ , whereas hyperbolicity of the pathological corner requires an active attenuation channel $\rho > 0$.

T3(b). Once $r_0 < \eta$, for $r_1 \geq a_1$ the drive $dr_1 \leq M - \kappa_1(a_1 - \eta) < 0$ carries r_1 below a_1 in finite time. Convergence to 0 then follows as below, the coupling $\kappa_1(r_0 - r_1)$ not being negative while $r_0 > r_1$. Fix $\delta > 0$. Since $r_0(t) \rightarrow 0$ there is T_δ with $r_0(t) < \delta$ for $t > T_\delta$; on the compact set $r_1 \in [\delta, a_1]$ the self term is uniformly negative while $\kappa_1(r_0 - r_1) \leq 0$, so dr_1 is uniformly negative and r_1 has no limit point there. Were $\limsup r_1 > 0$ it would have

one for small δ , so $r_1 \rightarrow 0$. Induction over $k = 1$ to n , with $r_{k-1} \rightarrow 0$ replacing $r_0 \rightarrow 0$, collapses the layer. Since $r_0 \rightarrow 0$, $z \rightarrow -a(Q_\infty) \leq -a_1 < 0$, so continuity of χ_- bounds the attenuation rate away from zero and $Q \rightarrow 0$. Necessity of some coupling is immediate from the $\kappa_k = 0$ dynamics; the bound is only the one used by the comparison argument.

T3(c1). On S , $z = 0$ gives $dr_0/dt = 0$, and $\chi_+(0) = \chi_-(0) = 0$ give $dQ/dt = 0$; hence $dz/dt = dr_0/dt - a'(Q) dQ/dt = 0$ and S is invariant, and by local Lipschitz continuity of the field (premise E) solutions are unique, so no trajectory leaves S or crosses it. For $z > 0$, $dQ/dt \geq 0$ gives $-a' dQ/dt \geq 0$, and $dr_0/dt = r_0(1-r_0)z > 0$, so $dz/dt > 0$; symmetrically $dz/dt < 0$ for $z < 0$. The root regime is therefore fixed by the sign of z at the excursion. By Proposition 1(a) the boundary conditions used here hold at every finite γ , so invariance does not require the sharp limit; the limit is what supplies exclusivity.

Proposition 2. On S , $z = 0$ gives $dr_0/dt = 0$, so $dz/dt = -a'(Q)dQ/dt$ and invariance requires $dQ/dt = 0$ at every point of S , that is $\alpha(0)C(r)(1 - Q) = \rho\beta(0)Q$ for all Q and all r with $r_0 = a(Q)$. The downstream coordinates are unconstrained on S , so $C(r)$ takes both zero and positive values there. Taking $r_1 = 0$ gives $\rho\beta(0)Q = 0$ for all Q , hence $\beta(0) = 0$ for $\rho > 0$; the remaining identity with $C(r) > 0$ and $Q < 1$ gives $\alpha(0) = 0$. Conversely $\alpha(0) = \beta(0) = 0$ gives $dQ/dt = 0$ on S and hence $dz/dt = 0$, so the condition is also sufficient within the class. Nothing forces the supports to be disjoint: exclusivity is

strictly stronger, not implied by invariance, and is what the one-sided evidence law of T4 and the monotonicity of z in T5 need. If $\alpha(0)$ and $\beta(0)$ are not both zero then $dz/dt \neq 0$ somewhere on S , so S is not invariant and the switching boundary acquires an ε dependence, as the identifiability-weighting row of Table 1 shows.

T3(c2). For $\Delta < \Delta_c(Q)$ the root is in $z > 0$, so r_0 rises to 1 (T3(c1)). Downstream, at a first touch $z_k = r_k - a(Q) = 0$ one has $dz_k/dt = \kappa_k(r_{k-1} - a(Q)) - a'(Q) dQ/dt > 0$, since $r_{k-1} > a(Q)$, true for $k = 1$ from the recovered start and then inductively, and $-a' dQ/dt \geq 0$; no r_k crosses its threshold, so $r_k > a(Q)$ for all t . Convergence to E_{rec} follows in three steps, the evidence first. The barrier gives $r_k > a(Q) \geq a_1$ for every k , so $C(r) \geq a_1^{n+1} > 0$ for all t , while $z(t) \geq z(0) > 0$ keeps the root in the recovery regime, so continuity of χ_+ gives $c_+ > 0$ with $\chi_+(z(t)) \geq c_+$; hence $dQ/dt \geq \varepsilon c_+ a_1^4 (1-Q)$ and $Q \rightarrow 1$. Next the root: $dr_0/dt = r_0(1-r_0)z > 0$ is monotone and bounded, so $r_0 \rightarrow 1$, a limit below 1 leaving $z > 0$ and a positive velocity. Then the chain, as a triangular scalar asymptotically autonomous cascade³⁷, as in the proof of T5 below, so $(r, Q) \rightarrow E_{\text{rec}}$. The recovered-configuration hypothesis is necessary: for $r_0 > a(Q)$ and every downstream node at 0 with small κ those nodes stay in the pathological basin, so no unconditional absorption can be claimed. For $\Delta > \Delta_c(Q)$ the root is in $z < 0$ and T3(b) applies.

Convergence proofs (Theorem 5)

This section gives in full the arguments summarised in T5; the numerical checks in the Supplementary Methods and Supplementary Fig. 1 are corroboration, not proof. Since $a \mapsto 1 - a + a^2$ is convex, $\max_x g_r(x, Q) = (1 - a(Q) + a(Q)^2)/3$ is maximised over $a(Q) \in [a_1, a_0]$ at an endpoint, so K^* is a uniform bound and $\kappa_k > K^*$ gives $g_r(x, Q) - \kappa_k < 0$ for every x and Q , each downstream field being then strictly decreasing in its own node; the bound is sufficient, not sharp. The regions $z > 0$ and $z < 0$ are forward invariant (T3(c1)), the field being subtangential to each³⁸. The two cases are proved separately rather than by symmetry, because the Q equation and the role of $C(r)$ differ: the recovery side needs a persistence estimate to keep the evidence rate away from zero, the collapse side does not.

Root convergence, recovery side. On $\{z > 0\}$ both r_0 and $-a(Q)$ are non-decreasing:

$dr_0/dt = r_0(1-r_0)z \geq 0$, and with $\chi_- = 0$ there $dQ/dt = \varepsilon\chi_+C(1-Q) \geq 0$ while $a' < 0$, so $dz/dt \geq 0$ and $z(t) \geq z(0) > 0$. The endpoint is separated: if $r_0(0) = 1$ then $r_0(t) \equiv 1$. Otherwise $r_0(0) \in (0, 1)$, since $r_0(0) > a(Q(0)) \geq a_1 > 0$, and $dr_0/dt \geq z(0) r_0(1-r_0)$; the comparison equation $dy/dt = z(0)y(1-y)$ with $y(0) = r_0(0)$ is logistic, so $y(t) \rightarrow 1$ monotonically and the comparison theorem gives $r_0(t) \geq y(t)$, which with $r_0 \leq 1$ forces $r_0(t) \rightarrow 1$.

Root convergence, collapse side. On $\{z < 0\}$, $\chi_+ = 0$, so $dQ/dt = -\varepsilon\rho\chi_-(z)Q \leq 0$ and,

since $a' < 0$, $-a'(Q)dQ/dt \leq 0$; with $dr_0/dt = r_0(1-r_0)z \leq 0$ this gives $dz/dt \leq 0$ and $z(t) \leq$

$z(0) < 0$. The endpoint is again separated: if $r_0(0) = 0$ then $r_0(t) \equiv 0$. Otherwise $r_0(0) \in (0,1)$, since $r_0(0) < a(Q(0)) \leq a_0 < 1$, and $dr_0/dt \leq z(0)r_0(1-r_0)$; comparison with $dy/dt = z(0)y(1-y)$, $y(0) = r_0(0)$, whose solution decreases monotonically to 0 because $z(0) < 0$, gives $r_0(t) \leq y(t) \rightarrow 0$, and $r_0 \geq 0$ forces $r_0(t) \rightarrow 0$. No persistence estimate is needed here, and none is available: the root falls under its own field, the coupling into it being absent ($\kappa_0 = 0$), and $C(r)$ is driven to zero rather than bounded below.

Persistence (recovery side). Suppose $\liminf r_{k-1}(t) \geq c_{k-1} > 0$, and let $\eta \in (0, c_{k-1})$, so that $r_{k-1}(t) \geq c_{k-1} - \eta$ for $t \geq \text{some } T$. Since $0 \leq r_k(1-r_k) \leq r_k$ and $r_k - a(Q) \geq -a_0$, one has $g(r_k, Q) \geq -a_0 r_k$, so $dr_k/dt \geq \kappa_k(c_{k-1} - \eta) - (a_0 + \kappa_k) r_k$ for $t \geq T$. The linear comparison equation has the globally attracting equilibrium $u^* = \kappa_k(c_{k-1} - \eta)/(a_0 + \kappa_k)$, so $\liminf r_k \geq u^*$, and $\eta \rightarrow 0$ gives $\liminf r_k \geq \kappa_k c_{k-1}/(a_0 + \kappa_k) > 0$. With $r_0 \rightarrow 1$ as base and inducting on $k = 1$ to n , $\liminf C(r) > 0$.

Evidence endpoint. On the recovery side Q is non-decreasing and bounded by 1, so $Q(t) \rightarrow Q_\infty$. Selector positivity gives $\chi_+(z(t)) \geq c_+ > 0$ once $z \geq z(0) > 0$, and $\liminf C(r) > 0$, so dQ/dt is bounded below by a positive multiple of $(1-Q)$ for large t , whence $Q_\infty = 1$. On the collapse side Q is non-increasing and bounded below, so $Q(t) \rightarrow Q_\infty$; $r_0 \rightarrow 0$ gives $z(t) \rightarrow -a(Q_\infty) \leq -a_1 < 0$, so z eventually lies in a compact subset of the active half-line of χ_- , where $\chi_-(z(t)) \geq c_- > 0$, whence $dQ/dt \leq -\epsilon pc_- Q$ and $Q_\infty = 0$.

Triangular cascade. The state chain is triangular: the equation for r_k depends only on r_k , on $r_{k-1}(t)$ and on $Q(t)$. Write $F_k(t, x) = g(x, Q(t)) + \kappa_k(r_{k-1}(t) - x)$ and, on the recovery side, $F_k^\infty(x) = g(x, 1) + \kappa_k(1 - x)$. Since g is uniformly continuous on the compact set $[0, 1]^2$ and $Q(t) \rightarrow 1$, $r_{k-1}(t) \rightarrow 1$, one has $\sup_{x \in [0, 1]} |F_k(t, x) - F_k^\infty(x)| \leq \sup_x |g(x, Q(t)) - g(x, 1)| + \kappa_k |r_{k-1}(t) - 1| \rightarrow 0$, so the field converges uniformly on $[0, 1]$ and the r_k equation is asymptotically autonomous in the sense the convergence theorem requires³⁷. Because $\kappa_k > K^*$, $(F_k^\infty)' = g_r(x, 1) - \kappa_k < 0$ on $[0, 1]$, so F_k^∞ is strictly decreasing with $F_k^\infty(1) = 0$, hence $x = 1$ is its only zero, positive below it and globally asymptotically stable on $[0, 1]$. The solution $r_k(t)$ is precompact in $[0, 1]$ and the limit semiflow has that single attracting equilibrium and no other compact invariant set, so $r_k(t) \rightarrow 1$. Inducting on $k = 1$ to n with base $r_0 \rightarrow 1$ gives $r \rightarrow (1, \dots, 1)$, and with $Q \rightarrow 1$ the trajectory converges to E_{rec} . On the collapse side the same argument applies with $r_{k-1}(t) \rightarrow 0$, $Q(t) \rightarrow 0$ and limit field $F_k^\infty(x) = g(x, 0) - \kappa_k x$, again strictly decreasing with unique zero $x = 0$, giving convergence to E_{path} . Because the nodes converge one stage at a time, no separate exclusion of periodic orbits is needed.

Dynamics on S and identification of the basins. On S the diagonal set $E_S = \{(a(Q), a(Q), a(Q), a(Q), Q)\}$ consists of equilibria, but a general point of S is not one, only the root velocity vanishing there. Since both selectors vanish on S , $Q(t) \equiv \bar{Q}$ is

frozen along any orbit in S , and the downstream chain is a triangular cascade of strictly decreasing fields with fixed input $r_0 \equiv a(\bar{Q})$, so $r_k(t) \rightarrow a(\bar{Q})$ for $k = 1$ to n by the same induction: every orbit within S converges to the point of E_S at its own evidence level. Transversally, $dz/dt = r_0(1-r_0)z - a'(Q)dQ/dt$ is positive for $z > 0$ and negative for $z < 0$, so E_S is repelling in the root direction and its points are non-attracting; the correct statement is two attractors off S , not two equilibria.

Under the strong-coupling condition the basins are identified exactly, by T5 rather than by continuity of z alone. Every trajectory in $\{z > 0\}$ converges to E_{rec} and every trajectory in $\{z < 0\}$ to E_{path} , so $B(E_{\text{rec}}) = \{z > 0\}$ and $B(E_{\text{path}}) = \{z < 0\}$; both are open, being preimages of open half-lines under the continuous map $(r, Q) \mapsto z$, and their union omits exactly S . Every point of S lies in the closure of both: there $a(Q) \in [a_1, a_0] \subset (0, 1)$, and for fixed downstream coordinates and Q the map $r_0 \mapsto z$ is strictly increasing with unit derivative, so displacing r_0 by $\pm\delta$ stays in the state cube and gives points of either sign arbitrarily close by. Hence $S \subseteq \partial B(E_{\text{rec}}) \cap \partial B(E_{\text{path}})$; neither open basin meets the boundary of the other, so $\partial B(E_{\text{rec}}) = \partial B(E_{\text{path}}) = S$, and S is exactly the common boundary of the two basins. The weak-coupling counterexample showing that the coupling condition cannot be dropped, the general linear evidence law, the proof of Proposition 3, the numerical verification and the scope of the premises are given in the

Supplementary Methods.

Use of generative artificial intelligence

A large language model (Claude, Anthropic) was used during manuscript preparation to assist with language editing, identification of potential logical gaps, and code review.

All mathematical arguments, proofs, numerical results and final wording were independently checked and approved by the author. The model was not used as an author and bears no responsibility for the work.

Data availability

This is a theoretical study; no datasets were generated or analysed.

Code availability

The reference implementation reproduces the results by status and regenerates every display item, Tables 1 to 3, Figures 1 and 2, and the supplementary figure. It is openly available at <https://github.com/entrance4-png/evidence-closed-loop> under the MIT licence and is archived under a persistent identifier at Zenodo³⁹.

References

1. Doignon, J.-P. & Falmagne, J.-C. Spaces for the assessment of knowledge. *Int. J. Man-Machine Studies* 23, 175-196 (1985).

2. Saito, H., Takizawa, Y., Tateno, A., Arakawa, R. & Tiger, M. Recovery from catatonia follows a hierarchical precision order. Preprint at medRxiv <https://doi.org/10.64898/2026.08.13.26360194> (2026).
3. Saito, H. The Saito Loop, a computational theory of hierarchical recovery in catatonia. Preprint at PsyArXiv <https://osf.io/preprints/psyarxiv/de6bu> (2026).
4. Saito, H. Recovery order in adaptive systems is set by dependency structure. Preprint at PsyArXiv <https://osf.io/preprints/psyarxiv/8ae7w> (2026).
5. Hirsch, M. W. Systems of differential equations that are competitive or cooperative II. *SIAM J. Math. Anal.* 16, 423-439 (1985).
6. Smith, H. L. *Monotone Dynamical Systems* (American Mathematical Society, 1995).
7. Allee, W. C. *Animal Aggregations: A Study in General Sociology* (Univ. Chicago Press, 1931).
8. Courchamp, F., Berec, L. & Gascoigne, J. *Allee Effects in Ecology and Conservation* (Oxford Univ. Press, 2008).
9. Scheffer, M., Carpenter, S., Foley, J. A., Folke, C. & Walker, B. Catastrophic shifts in ecosystems. *Nature* 413, 591-596 (2001).
10. Walther, S., Stegmayer, K., Wilson, J. E. & Heckers, S. Structure and neural mechanisms of catatonia. *Lancet Psychiatry* 6, 610-619 (2019).

11. Hirjak, D. et al. Catatonia. *Nat. Rev. Dis. Primers* 10, 49 (2024).
12. Rogers, J. P. et al. Evidence-based consensus guidelines for the management of catatonia: recommendations from the British Association for Psychopharmacology. *J. Psychopharmacol.* 37, 327-369 (2023).
13. Pelzer, A. C. M., van der Heijden, F. M. M. A. & den Boer, E. Systematic review of catatonia treatment. *Neuropsychiatr. Dis. Treat.* 14, 317-326 (2018).
14. Ali, S. F., Gowda, G. S., Jaisoorya, T. S. & Math, S. B. Resurgence of catatonia following tapering or stoppage of lorazepam: a case series and implications. *Asian J. Psychiatr.* 28, 102-105 (2017).
15. Lander, M., Bastiampillai, T. & Sareen, J. Review of withdrawal catatonia: what does this reveal about clozapine? *Transl. Psychiatry* 8, 139 (2018).
16. White, D. A. C. & Robins, A. H. Catatonia: harbinger of the neuroleptic malignant syndrome. *Br. J. Psychiatry* 158, 419-421 (1991).
17. Hirjak, D., Sartorius, A., Kubera, K. M. & Wolf, R. C. Antipsychotic-induced catatonia and neuroleptic malignant syndrome: the dark side of the moon. *Mol. Psychiatry* 26, 6112-6114 (2021).
18. Mann, S. C., Caroff, S. N., Bleier, H. R., Welz, W. K., Kling, M. A. & Hayashida, M. Lethal catatonia. *Am. J. Psychiatry* 143, 1374-1381 (1986).

19. Connell, J. et al. Malignant catatonia: a review for the intensivist. *J. Intensive Care Med.* 38, 137-150 (2023).
20. Ungvari, G. S., Caroff, S. N. & Gerevich, J. The catatonia conundrum: evidence of psychomotor phenomena as a symptom dimension in psychotic disorders. *Schizophr. Bull.* 36, 231-238 (2010).
21. Friston, K. The free-energy principle: a unified brain theory? *Nat. Rev. Neurosci.* 11, 127-138 (2010).
22. Rao, C. R. Information and the accuracy attainable in the estimation of statistical parameters. *Bull. Calcutta Math. Soc.* 37, 81-91 (1945).
23. Friston, K., Rigoli, F., Ognibene, D., Mathys, C., Fitzgerald, T. & Pezzulo, G. Active inference and epistemic value. *Cogn. Neurosci.* 6, 187-214 (2015).
24. Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P. & Pezzulo, G. Active inference: a process theory. *Neural Comput.* 29, 1-49 (2017).
25. Da Costa, L., Parr, T., Sajid, N., Veselic, S., Neacsu, V. & Friston, K. Active inference on discrete state-spaces: a synthesis. *J. Math. Psychol.* 99, 102447 (2020).
26. Mathys, C., Daunizeau, J., Friston, K. J. & Stephan, K. E. A Bayesian foundation for individual learning under uncertainty. *Front. Hum. Neurosci.* 5, 39 (2011).
27. Behrens, T. E. J., Woolrich, M. W., Walton, M. E. & Rushworth, M. F. S. Learning

the value of information in an uncertain world. *Nat. Neurosci.* 10, 1214-1221 (2007).

28. Adams, R. P. & MacKay, D. J. C. Bayesian online changepoint detection. Preprint at [arXiv:0710.3742](https://arxiv.org/abs/0710.3742) (2007).

29. Hoeting, J. A., Madigan, D., Raftery, A. E. & Volinsky, C. T. Bayesian model averaging: a tutorial. *Stat. Sci.* 14, 382-401 (1999).

30. Yu, A. J. & Dayan, P. Uncertainty, neuromodulation, and attention. *Neuron* 46, 681-692 (2005).

31. Nassar, M. R., Wilson, R. C., Heasly, B. & Gold, J. I. An approximately Bayesian delta-rule model explains the dynamics of belief updating in a changing environment. *J. Neurosci.* 30, 12366-12378 (2010).

32. Winters, C., van Wegen, E. E. H., Daffertshofer, A. & Kwakkel, G. Generalizability of the proportional recovery model for the upper extremity after an ischemic stroke. *Neurorehabil. Neural Repair* 29, 614-622 (2015).

33. Ogonek, J. et al. Immune reconstitution after allogeneic hematopoietic stem cell transplantation. *Front. Immunol.* 7, 507 (2016).

34. Kirkpatrick, J. et al. Overcoming catastrophic forgetting in neural networks. *Proc. Natl Acad. Sci. USA* 114, 3521-3526 (2017).

35. Bengio, Y., Louradour, J., Collobert, R. & Weston, J. Curriculum learning. In *Proc.*

26th International Conference on Machine Learning 41-48 (ACM, 2009).

36. Scheffer, M. et al. Early-warning signals for critical transitions. *Nature* 461, 53-59 (2009).

37. Thieme, H. R. Convergence results and a Poincare-Bendixson trichotomy for asymptotically autonomous differential equations. *J. Math. Biol.* 30, 755-763 (1992).

38. Nagumo, M. Über die Lage der Integralkurven gewöhnlicher Differentialgleichungen. *Proc. Phys.-Math. Soc. Japan* 24, 551-559 (1942).

39. Saito, H. evidence-closed-loop: reference implementation for "Evidence-coupled bistability converts dependency structure into a recovery geometry". Zenodo <https://doi.org/10.5281/zenodo.21669949> (2026).

Acknowledgements

Y. Takizawa, A. Tateno and R. Arakawa (Nippon Medical School and Onda-Daini Hospital), J. Theorell and M. Tiger (Karolinska Institutet) contributed to the companion clinical work. The theory presented here, and any errors in it, are the author's own.

Author contributions

H.S. is the sole author: he conceived the theory, developed the proofs, implemented the numerical verification and wrote the manuscript.

Competing interests

The author declares no competing interests.

Display items

Table 1. Structural dependence within the present construction.

Structure removed	Component lost
Bistability (Axiom B)	A distinct collapse basin
Concavity of $a(Q)$ (Axiom D)	Convex evidence-to-rate gain
Evidence gating $C(r)$ (Axiom C)	Evidence-mediated cross-level self-catalysis
Sharp regime selection ($\gamma \rightarrow \infty$ limit)	Regime exclusivity, hence the closed-form evidence law of T4 and the monotonicity of z used by T5; S itself survives at finite γ
Identifiability weighting at the boundary ($L(1/2) = 0$, E4)	Invariance of S at any sharpness, which Proposition 2 shows to be what invariance requires; the switching boundary becomes timescale dependent, drifting from 0.5940 at $\varepsilon = 0.02$ to 0.5786 at $\varepsilon = 0.00125$ against $Q_c = 0.5774$
Directional coupling (Axiom A)	Whole-layer propagation

Removing each listed structure eliminates at least the corresponding component. This does not assert logical independence or uniqueness of the assumptions.

Table 2. Status of theoretical results.

Result	Status
T0 positive invariance	Proved
T4 evidence law $Q = 1 - (1 - Q_0)\exp(-\varepsilon N)$, monotone in N	Proved on the forward-invariant recovery regime $\{z > 0\}$, given the premise; the attenuation counterpart holds on $\{z < 0\}$

	and the general linear solution is in the Supplementary Methods
T1 evidence-mediated acceleration, convexity, degree	Proved
T2 recovered state stable	Proved, unconditional in ρ
T2 pathological state stable	Proved for $\rho > 0$ (non-hyperbolic at $\rho = 0$)
T3(a) frozen-evidence root trigger	Proved (exact)
T3(b) whole-layer propagation	Proved under a sufficient coupling bound; some coupling necessary, bound not
T3(c1) exact invariant switching manifold of the root	Proved (all ε, ρ ; the boundary condition it uses holds at every finite γ)
T3(c2) whole-layer outcome and absorption	Proved from a recovered configuration (not over arbitrary data)
T5 two-outcome convergence off S, strong coupling	Proved for $\rho > 0$ and $\kappa_k > K^* \approx 0.291$, endpoint initial data included; K^* sufficient, not sharp; the range between the mixed-branch regime and K^* is not classified
T5 at finite sharpness (Proposition 6)	S invariant and the two sides exactly separated at every finite γ ; E_{path} exact; E_{rec} displaced to $1 - Q = \rho e^{-\gamma(1-a(Q))}Q$, within $\rho \exp[-\gamma(1 - a_0)]$ (Supplementary Methods)
Gate robustness (Proposition 7)	All results hold for any gate comparable to $\prod r_j$ with the same zero set; a leak leaves the theorems but removes the exact halting of accumulation (Supplementary Methods)
Exactness of S under a violation of E4	Boundary drift $O(\varepsilon\eta)$ for a boundary value $\chi_{\pm}(0) = \eta/2$; measured ratio 0.94 to 0.97 (Supplementary Methods)
Equilibrium continuum on S; stable mixed branch at sufficiently weak coupling	Characterised; $\kappa_1 < a_1^2/4$ gives a hyperbolically stable low first stage, $\kappa_1 = a_1^2/4$ the saddle-node (Supplementary Methods)
S is the global boundary of the two corner basins	Proved under T5 strong coupling, the basins being identified by T5 itself; not in the general T1 to T4 regime
E_{rec} and E_{path} are the only equilibria	False (equilibrium continuum E_S on S, to which orbits within S converge)
Proposition 1(a) boundary zeros and invariance of S at finite sharpness; 1(b)	Proved from E1 to E4 (Methods). The limit is discontinuous, so premise E

convergence to the indicator pair	regularises it and the theorems are proved on the regularisation
Proposition 2 necessity of boundary-zero updating for exact switching	Proved within the separable update class (Methods). Exclusivity is strictly stronger, needed for the one-sided laws of T4 and the monotonicity used in T5, not for invariance of S
Proposition 3 state order under root-initiated recovery	Proved (four state nodes; Supplementary Methods)
Proposition 4 closure: Q is the accumulated identifiability of the terminal node	Proved within the companion observation-model class, given the bridging conditions O1 to O3, which are what supply the exact evidence law; the first-level channel fixes the gate but not the rate. Proposition 4 derives the forward leg only; the return leg is derived conditionally from the separate decision-layer premises of Proposition 5 (Supplementary Methods)
Proposition 5: Axiom D from a decision rule	$a'(Q) < 0$ proved from expected-free-energy policy selection under the decision-layer premises D1 to D4; $a''(Q) < 0$ shown equivalent to $\Lambda^* \tilde{a}'' + \tilde{a}' < 0$, an explicit inequality, not a further premise (Supplementary Methods)
Milestone order of Q relative to the state nodes	Proved for sufficiently slow evidence under the closure reading (corollary to Proposition 4, Supplementary Methods); the five-parameter order is established in the companion theory under a different axiom class (refs 3, 4) and is not re-derived
Two-regime premise (E1), margin identification (E2), identifiability weighting (E4), regularisation (E)	Modelling commitments. E1 to E4 give the exclusive limit; E is its continuous one-sided regularisation, a well-posedness choice whose functional form no theorem uses
Clinical validity of the observation model	Empirical; out of scope here

All numerical values quoted in this table and in the text are illustrative and depend on the regularisation and parameter choice; the theorems depend only on the class properties of E.

Table 3. Division of labour across the four studies, and what does not transfer.

Study	Input	What it establishes	What it does not address
Companion clinical study (ref. 2)	Consecutive inpatients, order specified a priori	That the five capacities return in one order, without inversion	Why the order holds; any dynamics or endpoint
Companion computational study (ref. 3)	An active-inference architecture	The dependency graph, from the zero-pattern of the joint Fisher information; the chain has a unique linear extension, which is the observed order	The recovery dynamics and the outcomes it can reach
Graph-general recovery-order theory (ref. 4)	Any dependency graph, with availability gating, deficit-closing recovery and a common impaired onset	That the admissible orders are exactly the linear extensions of the graph, whatever the rate constants	Stability, consolidation, partial states and collapse
This paper	A forward acyclic dependency chain, carrying evidence-coupled bistable dynamics	The recovery geometry: evidence-mediated acceleration (T1), two stable configurations (T2), an exact invariant switching surface (T3), and the coupling condition separating a global dichotomy from a stable partial state (T5)	How the structure is obtained; the order under the companion gating axioms, none of which holds here, the order being recovered instead by forward coupling and root-initiated onset (Proposition 3); branching graphs with several parents per capacity




Figure 1. Dependency structure, evidence, and the outcomes it admits. **a**, Capacities r_0 to r_3 sit on a forward acyclic dependency chain with non-negative coupling κ_k , the root having no prerequisite. The evidence variable Q is not a node of that chain: it is generated by the capacities and acts back on all of them by lowering the shared threshold $a(Q)$, so the causal loop closes without introducing a backward edge. The catatonia instantiation identifies the chain with $\pi_s \rightarrow \beta \rightarrow \pi_m \rightarrow \pi_{v, \text{fast}} \rightarrow \pi_{v, \text{slow}}$. **b**, Evidence accrues at the rate $C(r) = \prod_j r_j$, which stays negligible until every prerequisite is online: dependency structure therefore gates when evidence can accumulate. Capacities are drawn recovering in the order the coupling enforces (Proposition 3). **c**, The bistable field $g(r, Q) = r(1-r)(r-a(Q))$ at three levels of accumulated evidence. Its interior zero is the unstable threshold $a(Q)$, which falls as evidence consolidates (Axiom D), widening the tolerated excursion $\Delta_c(Q) = 1 - a(Q)$ of T3(a) and accelerating every

capacity (T1). **d**, The three outcomes. Above the sufficient bound K^* every trajectory off the invariant surface $S = \{r_0 = a(Q)\}$ reaches full recovery or whole-system collapse according to the sign of $z(0)$ (T5); for sufficiently weak coupling, with $\kappa_1 < a_1^2/4$ at the first stage, a stable partial state exists even at fully consolidated evidence, the root restored and the dependent capacities near zero. Trajectories starting on S remain on it. Panels b to d are generated by the reference implementation at the illustrative parameters.




Figure 2. From the assumptions to the theorems. Axiom A, the forward dependency chain, splits into a recovery direction, in which evidence gating (Axiom C) yields the

evidence law (T4), concave basin deepening (Axiom D) then gives super-linear acceleration (T1) and a stable recovered state (T2); and a collapse direction, in which bistability and the regime-selective update (Axiom B and E), triggered by root autonomy, give the fragile window and whole-layer collapse (T3). The regime-selective update E is not an independent axiom: E1 to E4 give a Bayesian selector whose sharp-selection limit is exclusive (Proposition 1), and E is the continuous regularisation of that limit. Under strong coupling both directions resolve into the global two-outcome dichotomy (T5). The dashed edge is the self-catalytic feedback by which downstream recovery raises the evidence rate. The top row maps the paper: independently reported clinical features motivate the structural assumptions, which entail the theorems T1 to T5, which yield the testable predictions.

Supplementary Information

Evidence-coupled bistability converts dependency structure into a recovery geometry

Hiroki Saito

References are cited by the numbering of the main text.

Supplementary Methods

Weak coupling: the counterexample to T5 below threshold

For the counterexample we continue the mixed corner from $\kappa = 0$. At $\kappa = 0$ the point $(1,0,0,0,1)$ has each downstream node at a hyperbolic zero of $g(\cdot,1)$ with transverse eigenvalue $g_r(0,1) = -a_1 < 0$, so the implicit-function theorem continues it to a locally asymptotically stable mixed state equilibrium for all sufficiently small positive equal coupling. The continuation must be carried out on the invariant face $Q = 1$, not in the full system: at $\kappa = 0$ the factor $C = 0$ makes the Q -direction eigenvalue vanish, so the full five-dimensional Jacobian is singular in that direction and the implicit-function theorem does not apply there; on the face $Q = 1$ the four-node state subsystem is hyperbolic, so the theorem applies and the branch continues; on the continued branch positive κ gives every $r_k > 0$, hence $C(r) > 0$ and a full-system Q -eigenvalue

$-\epsilon\chi_+(1-a_1)C(r) < 0$, so the branch is a locally asymptotically stable equilibrium of the full system. The first-stage factor $r_1^2 - a_1r_1 + \kappa_1$ fixes the low roots of node 1 at $Q = 1$: for $\kappa_1 < a_1^2/4$ there are two distinct roots in $(0, a_1)$, of which the lower is hyperbolically stable in the r_1 direction; at $\kappa_1 = a_1^2/4$ the two coalesce at $r_1 = a_1/2$ in a saddle-node, where a low root exists but is not hyperbolic; and for $\kappa_1 > a_1^2/4$ node 1 has no low root. The strict inequality is therefore the condition for a hyperbolically stable low first stage and the non-strict one the condition for a low root to exist at all. Both are statements about that node, not about the existence of the whole branch, whose continuation is guaranteed only for sufficiently small κ .

Proof of Proposition 3

Set $d_k = r_{k-1} - r_k$ for $k = 1$ to n . With $u = 0$ the state equations give $dd_1/dt = g(r_0, Q) - g(r_1, Q) - \kappa_1 d_1$ and, for $k = 2$ to n , $dd_k/dt = g(r_{k-1}, Q) - g(r_k, Q) + \kappa_{k-1} d_{k-1} - \kappa_k d_k$. Consider the ordered cone $K = \{d_k \geq 0, k = 1 \text{ to } n\}$, which contains the initial data. On the face $d_k = 0$ one has $r_{k-1} = r_k$, so the self terms cancel and the term $-\kappa_k d_k$ vanishes; hence $dd_1/dt = 0$ there for $k = 1$, and $dd_k/dt = \kappa_{k-1} d_{k-1} \geq 0$ there for $k = 2$ to n , the last inequality holding because the point lies in K . The field is therefore subtangential to K at every boundary point, and the right-hand side is locally Lipschitz on the compact state space, so K is forward invariant³⁸. The argument uses no property of $a(Q)$ beyond continuity,

so it holds in the recovery and the collapse regime alike and for every ε and p . Since $r_{k-1}(t) \geq r_k(t)$ for all t , the first time r_{k-1} attains a common threshold q is no later than the first time r_k does, which is the stated ordering of first-passage times, read in the extended reals with $T_k(q) = \infty$ when the threshold is never attained.

Two boundaries. First, the proposition orders the four state variables and not the evidence coordinate: $dQ/dt = \varepsilon \chi_+(z) C(r)(1 - Q)$ is positive whenever the states are positive and the root is in the recovery regime, so for suitable $Q(0)$ and ε the evidence can cross a numerical threshold before r_3 does, and ordering Q against the states needs a separate milestone convention or a condition on ε , neither of which is assumed here; the corollary to Proposition 4 below supplies such a condition in the catatonia instantiation. Second, the hypothesis is an ordered root-initiated configuration and not merely $z(0) > 0$. In the unforced system the state diagonal $\{r_0 = r_1 = r_2 = r_3\}$ is invariant, because the diffusive terms vanish there and every node then follows the same field $g(r, Q)$, so an exactly uniform onset recovers synchronously; ordered recovery requires the symmetry breaking supplied by $u(t)$. This is also what places the present dynamics outside the companion recovery-order axioms, which assume a common impaired onset.

Proofs of T1, T3(a) and T4

T4. On the recovery regime $\chi_- = 0$, so $dQ/dt = \varepsilon \chi_+ C(1 - Q)$; with $w = 1 - Q$ this is

$dw/dt = -\varepsilon \chi_+ C w$, giving $w(t) = (1-Q_0)\exp[-\varepsilon N(t)]$ and the stated form, of which $Q_0 = 0$ is a special case. The monotonicity is in N , not in time: where any capacity is offline $C(r) = 0$, so $dN/dt = 0$ and N and Q are momentarily constant. The law $dQ = \varepsilon C(1-Q)$ is the exact differential of $Q = 1 - \exp[-(\Lambda - \Lambda_0)/\Lambda^*]$ under $d\Lambda/dt = I_0 C$ (Axiom C), with $\varepsilon = I_0/\Lambda^*$: a derivation, not a posit. On $\{z < 0\}$ the same computation with $\chi_+ = 0$ gives the attenuation form.

General evidence law. Write $A(t) = \varepsilon \chi_+(z(t))C(r(t)) \geq 0$ and $B(t) = \varepsilon \rho \chi_-(z(t)) \geq 0$, both continuous and bounded. The evidence equation $dQ/dt = A(t) - [A(t) + B(t)]Q$ is linear, so the integrating factor $\exp \int_0^t (A + B)$ gives $Q(t) = e^{-\int_0^t (A+B)} [Q_0 + \int_0^t A(s) e^{\int_0^s (A+B)} ds]$, preserving $0 \leq Q \leq 1$. This reduces to the closed form of T4 when $B \equiv 0$ and to pure attenuation when $A \equiv 0$. Under E and T3(c1) every trajectory satisfies one of the two, so T4 and its counterpart cover all trajectories of the sharp system; the general form is what finite γ needs, both channels then acting on every trajectory.

T1. Direct differentiation of g and C , using $a' < 0$, $a'' < 0$ (Axiom D) and $0 < r_k < 1$. The first inequality shows that Q raises each capacity's rate, the route being $r_{\text{downstream}} \rightarrow C(r) \rightarrow Q \rightarrow dr_{\text{upstream}}/dt$; the second is super-linearity, which fails when a is linear; the third is complementarity. What is proved is convexity of the map from Q to dr_k/dt , not time-convexity of the trajectory.

Corollary (degree). With $a(Q) = a_0 - (a_0 - a_1) Q^p$, the relative curvature is $a''(Q)/a'(Q) = (p-1)/Q$, independent of r , a_0 and a_1 ; zero at $p = 1$, positive for $p > 1$, and largest at low Q because $a'(Q)$ itself vanishes as $Q \rightarrow 0$. The absolute curvature $p(p-1)(a_0 - a_1) r_k(1 - r_k) Q^{p-2}$ is constant in Q for $p = 2$, so the low- Q dominance is relative. The complementarity elasticity is $1/r_j$.

T3(a). $dr_0/dt = g(r_0, Q)$ has the single interior unstable zero $r_0 = a(Q)$, with $r_0 < a(Q)$ giving $r_0 \rightarrow 0$ and $r_0 > a(Q)$ giving $r_0 \rightarrow 1$. The criterion is exact because the root has no incoming coupling, and $\Delta_c(Q) = 1 - a(Q)$ is increasing in Q by Axiom D.

Closure of the loop in the catatonia instantiation

This section derives, rather than assumes, the observation-model premise of the Results for the catatonia instantiation. What has to be shown is that the statistic which identifies the terminal node of the companion chain and the evidence variable Q of this paper are the same object.

The companion observation model. The companion computational study specifies a joint probabilistic observation model for $\theta = (\pi_s, \beta, \pi_m, \pi_{v_{\text{fast}}}, \pi_{v_{\text{slow}}})^3$: five conditionally independent Poisson channels with means $\mu_j(\theta) = M_j(\theta) \exp(\theta_j)$, where $\phi(x) = 1 - e^{-x}$, the enactment probability is $e = 1 - \exp(-a_E \beta \phi(\pi_s))$ with $a_E > 0$, and the sample-count functions are $M_s \propto 1$, $M_\beta \propto \phi(\pi_s)$, $M_m \propto e \phi(\pi_s)$, $M_f \propto e \phi(\pi_m)$ and M_{slow}

$\propto e^{\varphi(\pi m) - \varphi(\pi v_{\text{fast}})}$, these being the commitments M1 to M3. Its Lemma 4 and the mean-Jacobian column test give that the slow channel is the only one whose mean depends on πv_{slow} .

Bridging conditions. Three further conditions connect that model to the dynamics used here. They are commitments of the observation model, of the same status as M1 to M3, not consequences of it. **O1 (availability reading).** $r_k(t)$ in $[0,1]$ is the probability that node k is online at time t , and the four nodes are conditionally independent given the capacities, so the probability that all are online at once is $\prod_k r_k = C(r)$. **O2 (context sampling).** Contexts arrive at a rate v that does not depend on θ ; a context yields a usable per-context estimate of πv_{fast} if and only if every node is online throughout it; and the capacities vary slowly on the scale of one context. **O3 (per-estimate information).**

In the second level of the hierarchical model each usable per-context estimate contributes the same Fisher information $\iota > 0$ about πv_{slow} , while a context in which any node was offline leaves the likelihood exactly flat in πv_{slow} and contributes none.

Lemma C1 (the slow count is a product gate). Put $r_0 = \varphi(\pi s)$, $r_1 = \varphi(a_E \beta)$, $r_2 = \varphi(\pi m)$ and $r_3 = \varphi(\pi v_{\text{fast}})$, each in $[0,1)$, and let M^* be the normalising count of the slow channel. If $a_E \beta \leq B$ then $C(r) \leq M_{\text{slow}}/M^* \leq K C(r)$ with $K = B/(1 - e^{-B}) \geq 1$. Proof.

Write $c = a_E \beta$ in $[0,B]$. The map $u \mapsto 1 - e^{-cu}$ is concave on $[0,1]$ and vanishes at $u = 0$,

so it lies above the chord joining its endpoints and below its tangent at the origin: $r_0(1 - e^{-c}) \leq 1 - \exp(-c r_0) \leq c r_0$. The left member is $r_0 r_1$, so $c \geq r_0 r_1$. For the upper bound, concavity of $c \mapsto 1 - e^{-c}$ on $[0, B]$ gives $1 - e^{-c} \geq c(1 - e^{-B})/B$, that is $c \leq K r_1$, whence $c \leq K r_0 r_1$. Multiplying by $\varphi(\pi m)\varphi(\pi v_{\text{fast}}) = r_2 r_3$ gives the statement. In particular M_{slow} vanishes exactly when some r_k vanishes, which is the zero set Axiom C requires, and the two gates agree up to a factor in $[1, K]$.

Lemma C2 (accumulated information, first-level reading). For a Poisson channel with mean $\mu(\theta) = M e^\theta$ one has $\partial\mu/\partial\theta = \mu$, so its Fisher information²² about θ is $(\partial\mu/\partial\theta)^2/\mu = \mu$. Since no other channel depends on πv_{slow} , the joint Fisher information about it is $\mu_{\text{slow}} = M_{\text{slow}} \exp(\pi v_{\text{slow}})$, and information is additive over conditionally independent sampling epochs, so $\Lambda(t) - \Lambda_0 = \int_0^t \mu_{\text{slow}}(s) ds$. By Lemma C1 this lies between $M^* \exp(\pi v_{\text{slow}}) \int_0^t C(r) ds$ and K times that quantity. The prefactor $\exp(\pi v_{\text{slow}})$ is not a constant of the motion: πv_{slow} is itself recovering along the clinical chain, so this reading fixes the zero set and the dependence on $C(r)$ but not the rate.

Lemma C3 (accumulated information, hierarchical reading). Supplementary Table 1 of the companion study records that πv_{slow} has no first-level statistic of its own: its statistic is the set of per-context estimates of πv_{fast} . Under O1 to O3 the expected number of usable estimates in $[0, t]$ is $v \int_0^t C(r(s)) ds$, each contributing ι and the

unusable contexts contributing nothing, so $\Lambda(t) - \Lambda_0 = \imath \vee \int_0^t C(r(s)) \, ds$.

Proposition 4 (closure of the loop). Under M1 to M3 and O1 to O3, Axiom C holds in the exact form $d\Lambda/dt = I_0 C(r)$ with $I_0 = \imath \vee$, and $Q(t) = 1 - \exp[-(\Lambda(t) - \Lambda_0)/\Lambda^*]$ is a strictly increasing bijection from the Fisher information about $\pi_{V_{\text{slow}}}$ accumulated by time t onto $[0,1)$, with $\varepsilon = I_0/\Lambda^*$. The evidence variable of this paper and the statistic that identifies the terminal node of the companion chain are therefore the same object, read on a bounded scale. Proof. Lemma C3 is Axiom C with $I_0 = \imath \vee$, and it is the hierarchical reading alone that supplies the law in that exact form. Lemma C2 agrees with it on the zero set and on the dependence on $C(r)$, by the bracket of Lemma C1, but its prefactor $M^* \exp(\pi_{V_{\text{slow}}})$ varies as $\pi_{V_{\text{slow}}}$ recovers, so the first-level reading corroborates the gate structure rather than fixing the rate. The map $\Lambda \mapsto 1 - \exp[-(\Lambda - \Lambda_0)/\Lambda^*]$ is strictly increasing from $[\Lambda_0, \infty)$ onto $[0,1)$, and T4 is its differential form.

What the closure settles and what it does not. The forward leg is derived: within this observation model the joint availability of the four capacities is the only source of information about $\pi_{V_{\text{slow}}}$, and Q is that information on a bounded scale, so the near-identity the two papers seemed to share is an identity. The return leg is not derived here. That accumulated slow-volatility precision lowers the threshold shared by the capacities is Axiom D; nothing in the observation model forces it, and this closure concerns the

identity of the two quantities, not the sign of the feedback. The sign is settled separately, by the decision layer of Proposition 5 below. Both readings are conditional on the class of observation models sharing the structural zero set of M1 to M3, exactly as the companion uniqueness claims are, and any member of that class gives the same product gate; the exact evidence law rests on the hierarchical reading, that is on O1 to O3.

Outside catatonia the identification of Q with an accumulated evidence remains the premise stated in the Results.

Corollary (the evidence coordinate completes last). Assume $\kappa_k > K^*$, $r_0(0) > a_0$ and $Q(0) = Q_0$. Then for every q in $(Q_0, 1)$ there is $\varepsilon_0 > 0$ such that for $0 < \varepsilon < \varepsilon_0$ every capacity reaches q strictly before Q does. Proof. Since $a(Q) \leq a_0$, the field dominates componentwise that of the comparison system $d\tilde{r}_k/dt = \tilde{r}_k(1 - \tilde{r}_k)(\tilde{r}_k - a_0) + \kappa_k(\tilde{r}_{k-1} - \tilde{r}_k)$ with the same initial data, which is cooperative, so $r_k(t) \geq \tilde{r}_k(t)$ for all t and k ⁶. In the comparison system $\tilde{r}_0(0) > a_0$ gives $\tilde{r}_0 \rightarrow 1$, and $\kappa_k > K^* \geq \max_x \partial_t[x(1-x)(x-a_0)]$ makes each downstream field strictly decreasing, so the triangular cascade argument of T5 gives $\tilde{r}_k \rightarrow 1$ for every k ; hence the first-passage times $\tilde{T}_k(q)$ are finite and do not depend on ε , and $T_k(q) \leq \tilde{T}_k(q)$. On the other hand $\chi_+ \leq 1$ and $C \leq 1$ give $N(t) \leq t$, so $Q(t) \leq 1 - (1 - Q_0)e^{-\varepsilon t}$ and $T_Q(q) \geq \varepsilon^{-1} \log[(1 - Q_0)/(1 - q)]$. Choosing ε_0 so that the right-hand side exceeds $\max_k \tilde{T}_k(q)$ proves the claim. This supplies the condition on ε left

open in the two boundaries above: under the identifiability reading of Proposition 4 the evidence coordinate is slow exactly when the information required, Λ^* , is large compared with what one context supplies, and the slow-volatility milestone is then the last to be reached, which is the reported clinical order.

Axiom D from an active-inference decision rule

Axiom D asserts two signs, $a'(Q) < 0$ and $a''(Q) < 0$. Neither is a posit once the rule by which a policy is committed to is written down. None of the companion studies contains such a rule, so a decision layer is added here in the standard active-inference form²³⁻²⁵.

It is an additional premise set, D1 to D4, not a consequence of Axioms A to C.

D1 (two policies at each capacity). At capacity k the agent compares enact, which exercises the capacity, with withhold, which does not. Policies are selected by expected free energy, the posterior being $\sigma(-\gamma G)$ with $\gamma > 0$ the policy precision^{24,25}, so enactment occurs when $G(\text{enact}) < G(\text{withhold})$, and in the sharp limit of that softmax the comparison is a threshold rule, as in premise E.

D2 (generative model of the outcome). Under enact the predicted outcome is Gaussian, $o \sim N(\mu(r), v)$, with μ strictly increasing in the current availability r of the capacity and $\mu(r) < m$ on $[0,1)$, and $v = \zeta^2 + 1/\Lambda$, where ζ^2 is the irreducible outcome noise and $1/\Lambda$ the posterior variance of the slow environmental model whose

accumulated Fisher information is Λ . Prior preferences are Gaussian with mean m and precision λ_P . This is the standard linear-Gaussian propagation: uncertainty about the slow model inflates the predicted spread of the consequences of acting, which is the sense in which volatility precision is the quantity at stake^{26,27}.

D3 (the withhold branch does not use the slow model). Withholding produces the default outcome, whose prediction does not depend on the slow model, so $G(\text{withhold}) = G_0$ is constant in Λ and in r .

D4 (use-dependent maintenance). A capacity grows when it is exercised and decays when it is not, within the bounded envelope $r(1 - r)$, so the sign of dr_k/dt agrees with the enactment decision at that capacity. The unstable threshold of Axiom B is therefore the commitment threshold of D1.

Expected free energy decomposes into risk and ambiguity^{23,25}, $G = D_{KL}[Q(o|\pi) \parallel P(o)] + E[H[P(o|s, \pi)]]$, so for the model of D2

$$G_1(r, \Lambda) = \frac{1}{2}[\lambda_P v + \lambda_P(\mu(r) - m)^2 - 1 - \log(\lambda_P v)] + \frac{1}{2} \log(2\pi e \zeta^2), \quad v = \zeta^2 + 1/\Lambda,$$

the ambiguity term being independent of r and of Λ .

The eight hypotheses. (1) Enact and withhold are compared by expected free energy.

(2) Under enact the outcome is $o \sim N(\mu(r), \zeta^2 + 1/\Lambda)$. (3) $\mu'(r) > 0$ and $\mu(r) < m$, so a

higher capacity moves the predicted outcome towards the preferred one without

reaching it at $r < 1$. (4) Preferences are Gaussian with precision λ_P . (5) The expected free energy of withhold does not depend on r or Λ . (6) The informative regime $\lambda_P v > 1$ holds. (7) The threshold is interior, $G_1(1, \Lambda) < G_0 < G_1(0, \Lambda)$. (8) The sign of the capacity field agrees with the policy decision. Hypotheses (1) to (5) and (8) are D1 to D4 written out; (6) and (7) are regime conditions.

Proposition 5 (Axiom D from the decision rule). Assume D1 to D4, the informative regime $\lambda_P v > 1$, and the non-degenerate case $G_1(1, \Lambda) < G_0 < G_1(0, \Lambda)$. Then (a) the enactment set is the interval $\{r > a\}$ with $a = \tilde{a}(\Lambda)$ the unique solution of $G_1(a, \Lambda) = G_0$, and $\tilde{a}'(\Lambda) < 0$: sharpening the slow model lowers the threshold for committing to the policy. (b) With Q the evidence coordinate of the Results, $a''(Q) < 0$ if and only if $\Lambda^* \tilde{a}''(\Lambda) + \tilde{a}'(\Lambda) < 0$, equivalently if and only if $\Lambda > q(\Lambda) \Lambda^*$, where $q(\Lambda) := \Lambda \tilde{a}''(\Lambda) / |\tilde{a}'(\Lambda)|$ is the elasticity of the threshold's sensitivity to precision.

Proof. (a) $\partial G_1 / \partial r = \lambda_P (\mu(r) - m) \mu'(r) < 0$, since $\mu' > 0$ and $\mu(r) < m$, so G_1 is strictly decreasing in r and $\{G_1 < G_0\}$ is an upper interval whose endpoint is a , unique and interior by the non-degeneracy assumption. Next $\partial G_1 / \partial \Lambda = \frac{1}{2}(\lambda_P - 1/v) dv/d\Lambda$ with $dv/d\Lambda = -1/\Lambda^2 < 0$, so in the informative regime $\partial G_1 / \partial \Lambda < 0$. Since $\partial G_1 / \partial r \neq 0$ the implicit function theorem applies and $\tilde{a}'(\Lambda) = -(\partial G_1 / \partial \Lambda) / (\partial G_1 / \partial r) < 0$. (b) Inverting $Q = 1 - \exp[-(\Lambda - \Lambda_0)/\Lambda^*]$ gives $\Lambda(Q) = \Lambda_0 - \Lambda^* \log(1 - Q)$, with $\Lambda' = \Lambda^* / (1 - Q)$ and Λ''

$= \Lambda^*/(1 - Q)^2$. Hence $a''(Q) = \tilde{a}''(\Lambda)\Lambda'^2 + \tilde{a}'(\Lambda)\Lambda'' = [\Lambda^*/(1 - Q)^2][\Lambda^*\tilde{a}''(\Lambda) + \tilde{a}'(\Lambda)]$,

and the prefactor is positive.

Reading of the two parts. Part (a) is the statement that higher precision of the slow model lowers the threshold for committing to the current policy, and it is a theorem of the decision rule rather than an assumption: the only substantive condition is the informative regime $\lambda_P v > 1$, which says that the predictive spread still exceeds the tolerance of the preferences, so there is something left for precision to buy. Part (b) is an equivalence and not a further premise: concavity in Q says exactly that the evidence coordinate does not saturate faster than the marginal benefit of precision decays. If $|\tilde{a}'|$ falls like Λ^{-q} with q constant the condition reads $\Lambda > q\Lambda^*$, so it holds throughout when the baseline information Λ_0 already exceeds $q\Lambda^*$.

A worked instance. Take $\mu(r) = m - b(1 - r)$, so that $G_1(a, \Lambda) = G_0$ gives $\frac{1}{2}\lambda_P b^2(1 - a)^2 = F(\Lambda)$ with $F(\Lambda) = G_0 + \frac{1}{2} - \frac{1}{2}\log(2\pi e\zeta^2) + \frac{1}{2}\log(\lambda_P v) - \frac{1}{2}\lambda_P v$, and $a(\Lambda) = 1 - (2F(\Lambda)/\lambda_P b^2)^{1/2}$. Differentiating, $F'(\Lambda) = \frac{1}{2}(1/v - \lambda_P)v'(\Lambda) > 0$ in the informative regime, since $1/v - \lambda_P < 0$ and $v' < 0$, and $\tilde{a}'(\Lambda) = -F'(\Lambda)/(2\lambda_P b^2 F(\Lambda))^{1/2} < 0$, so the first sign is visible in closed form as well as through the implicit function theorem. Two constants fix the endpoints, $a(\Lambda_0) = a_0$ and $a(\Lambda) \rightarrow a_1$ as $\Lambda \rightarrow \infty$. At $\lambda_P = 2$, $\zeta^2 = 1$, $\Lambda_0 = 10$ and $\Lambda^* = 2$, with $a_0 = 0.60$ and $a_1 = 0.15$, the derived threshold is strictly decreasing and strictly

concave on the whole of $[0,1)$: the elasticity is $q(\Lambda_0) = 3.91$, falling to 2 as $\Lambda \rightarrow \infty$, and $\max_{\Lambda} q(\Lambda)\Lambda^*/\Lambda = 0.78 < 1$, which is the condition of part (b). Raising the coordinate scale to $\Lambda^* = 5$, so that the same maximum becomes $1.95 > 1$, leaves part (a) intact and makes a convex over part of the range, exactly as the equivalence predicts. The realisation used in the Results, $a(Q) = a_0 - (a_0 - a_1)Q^p$, is one admissible member of the same family, concave for $p > 1$ and borderline at $p = 1$; no theorem uses either form.

What remains a commitment. D1 to D4 are premises of the decision layer, of the same status as M1 to M3 of the observation model, and are not consequences of active inference in general: the reduction to two policies, the Gaussian outcome model, the constancy of the withhold branch and use-dependent maintenance are each modelling choices. What is derived is that, once such a layer is written down, the direction of the feedback is fixed and its curvature is settled by one explicit inequality. If the ambiguity term is evaluated under the posterior over the slow parameter rather than at its mode it contributes a further negative term to $\partial G_1/\partial \Lambda$, so part (a) is robust to that refinement. Together with Proposition 4 this closes both legs. The return leg is the chain $Q \uparrow \Rightarrow \Lambda \uparrow \Rightarrow G_1 \downarrow \Rightarrow a(Q) \downarrow \Rightarrow dr/dt \uparrow$, so the self-reinforcing path $r \rightarrow C(r) \rightarrow \Lambda \rightarrow Q \rightarrow a(Q) \rightarrow dr/dt$ is derived under the stated observation and decision models rather than assumed. The precise status is that $a'(Q) < 0$ is derived from the decision rule, $a''(Q) < 0$

is derived under an explicit inequality, and the two-policy reduction, the Gaussian outcome model, the constant withhold branch and use-dependent maintenance are the primitive premises: Axiom D is derived conditionally from a specified active-inference decision layer.

T5 at finite sharpness

T5 is proved for selectors in the class E, whose supports are disjoint, whereas the inferentially defensible object is the finite- γ pair of Proposition 1(a), in which both channels act everywhere. This section states what survives, and it is more than an approximation.

What survives exactly. First, $\chi_{\pm}(0) = 0$ holds at every finite γ , by Proposition 1(a), so by Proposition 2 the surface S is invariant and, by local Lipschitz uniqueness, $\{z > 0\}$ and $\{z < 0\}$ are forward invariant. The sign of z is therefore a conserved label at every sharpness and the classification of initial data by that sign is exactly the one T5 uses.

Second, E_{path} is an exact equilibrium at every finite γ : at $r = 0$ the gate $C(r)$ vanishes and at $Q = 0$ the attenuation term vanishes, so $dQ/dt = 0$ there, and the state block is unchanged.

What moves. On the invariant face $r = (1, \dots, 1)$ the state equations vanish identically and, since $\chi_-/\chi_+ = e^{-\gamma z}$ exactly for the E4 class, the evidence obeys $dQ/dt = \varepsilon \chi_+(z)[(1 - Q)$

$- \rho e^{-\gamma z} Q]$ with $z = 1 - a(Q)$. At $Q = 1$ this is strictly negative, so E_{rec} is not an equilibrium at finite γ . Its zeros satisfy $1 - Q = \rho e^{-\gamma(1-a(Q))} Q$, and since $a(Q) \leq a_0$ every zero obeys $1 - Q_\gamma \leq \rho \exp[-\gamma(1 - a_0)]$; once that holds, $a(Q_\gamma)$ is close to a_1 and the sharper estimate $1 - Q_\gamma \leq \rho \exp[-\gamma(1 - a_1)](1 + o(1))$ follows.

Lemma F (uniform margin at finite sharpness). Let the E4 member satisfy $L(w) \leq C_L(w - 1/2)^2$ for $|w - 1/2| \leq 1/4$; the normalised entropy deficit does so with $C_L = 3.1$. Put $A_1 = \max|a'|$, $M_\rho = \max(1, \rho)$ and $m^* = \min\{(a_1/2)(1 - a_1/2), (1 + a_0)(1 - a_0)/4\}$, a constant of the parameters alone, and set $\bar{z} = \min\{a_1/2, (1 - a_0)/2, 1/\gamma, 8 m^*/(\varepsilon A_1 M_\rho C_L \gamma^2)\} > 0$. Then $dz/dt > 0$ whenever $0 < z \leq \bar{z}$ and $dz/dt < 0$ whenever $-\bar{z} \leq z < 0$, so $|z(t)| \geq \zeta := \min(|z(0)|, \bar{z}) > 0$ for all $t \geq 0$.

Proof. $dz/dt = r_0(1 - r_0)z - a'(Q)dQ/dt$ with $|dQ/dt| \leq \varepsilon M_\rho \max(\chi_+, \chi_-)$ and $|a'| \leq A_1$. The constants are not circular and the same estimate covers both signs of z : the first two bounds in $\bar{z}$ are imposed before m^* is used, so whenever $0 < |z| \leq \bar{z}$ one has $r_0 = a(Q) + z$ in $[a_1/2, (1 + a_0)/2] \subset (0, 1)$, because $a(Q) \in [a_1, a_0]$ and $|z| \leq \min\{a_1/2, (1 - a_0)/2\}$; since $x(1 - x)$ is minimised on that interval at an endpoint, $r_0(1 - r_0) \geq m^*$. The third bound gives $|w_+ - 1/2| = \frac{1}{2}|\tanh(\gamma z/2)| \leq \gamma|z|/4 \leq 1/4$, so the quadratic bound on L applies and both $\chi_\pm \leq L(w_+) \leq C_L \gamma^2 z^2/16$. Hence $z \cdot dz/dt \geq z^2[m^* - \varepsilon A_1 M_\rho C_L \gamma^2 |z|/16]$, and the fourth bound in $\bar{z}$ makes the bracket at least $m^*/2 > 0$, so $dz/dt \geq m^*z/2 > 0$ for $0 < z \leq \bar{z}$ and

$dz/dt \leq m^*z/2 < 0$ for $-\bar{z} \leq z < 0$. Were $|z(t_1)| < \zeta$ for some t_1 there would be a first time at which $|z|$ equals $\zeta \leq \bar{z}$ and $|z|$ is not increasing, contradicting this.

The lemma replaces the monotonicity of z used by T5: at finite γ the margin z is no longer monotone but it is bounded away from the switching surface by an explicit amount. Members whose L vanishes only linearly at $1/2$, such as $|2w - 1|$, give $\chi_{\pm} \leq \gamma z/4$ and the same conclusion under the weaker requirement $\varepsilon A_1 M_p C_L \gamma < 2m^*$.

Proposition 6 (dichotomy at finite sharpness). Let $\gamma < \infty$, $\rho > 0$, $\kappa_k > K^*$, and let the selector be the finite-sharpness pair of Proposition 1(a) with L as in Lemma F. Then (a) S is invariant and the two regions are forward invariant, so the classification by the sign of $z(0)$ is exactly that of T5; (b) if $z(0) < 0$ then $(r, Q) \rightarrow E_{\text{path}}$, which is an exact equilibrium; (c) if $z(0) > 0$ then $r(t) \rightarrow (1, \dots, 1)$ and the ω -limit set lies in $\{(1, \dots, 1, Q) : 1 - Q = \rho e^{-\gamma(1-a(Q))Q}\}$, hence within $\rho \exp[-\gamma(1 - a_0)]$ of E_{rec} . The capacity-level dichotomy therefore holds exactly at every finite sharpness, the recovered evidence coordinate alone being displaced by an amount exponentially small in γ , and the exclusive limit recovers T5 as $\gamma \rightarrow \infty$.

Proof. (a) is Proposition 1(a) with Proposition 2 and uniqueness. For (b), Lemma F gives $z(t) \leq -\zeta < 0$, so $dr_0/dt \leq -\zeta r_0(1 - r_0)$ and the logistic comparison gives $r_0 \rightarrow 0$.

The chain then collapses under the strong-coupling condition of the proposition itself,

not under the different bound of T3(b): since $g(0, Q) = 0$ and $g_r(x, Q) \leq K^*$ uniformly, the mean value theorem gives $g(x, Q) \leq K^*x$ on $[0, 1]$, hence $dr_k/dt \leq \kappa_k r_{k-1} - (\kappa_k - K^*)r_k$ with $\kappa_k - K^* > 0$. Scalar linear comparison with the input $r_0 \rightarrow 0$ gives $r_1 \rightarrow 0$, and induction on k gives $r_k \rightarrow 0$ for every k , so $C(r) \rightarrow 0$. Since $z \leq -\zeta$ keeps $\chi_-(z(t)) \geq c_- > 0$ on the compact range of z , the evidence equation gives $dQ/dt \leq \varepsilon \chi_+ C(1 - Q) - \varepsilon p c_- Q$, whose first term tends to zero, so $Q \rightarrow 0$. For (c), Lemma F gives $z(t) \geq \zeta > 0$, so $dr_0/dt \geq \zeta r_0(1 - r_0)$ and $r_0 \rightarrow 1$; the persistence estimate of T5 is unchanged and gives $\liminf C(r) > 0$, so with $\chi_+(z) \geq c_+ > 0$ and $e^{-\gamma z} \leq e^{-\gamma \zeta}$ the evidence satisfies $dQ/dt \geq \varepsilon c_+ [c_c(1 - Q) - \rho e^{-\gamma \zeta} Q]$ and is eventually above a level $\bar{Q} < 1$ with $a(\bar{Q}) < 1$. Comparison with the frozen-threshold chain at $a \equiv a(\bar{Q})$, which is cooperative and triangular with $\kappa_k > K^*$, gives $r_k \rightarrow 1$ for every k . The ω -limit set is therefore contained in the invariant face $r = (1, \dots, 1)$, on which the flow is scalar in Q , so it consists of zeros of that scalar field, which are the stated points.

Robustness of the evidence gate

Proposition 7 (gates comparable to the product). Replace $C(r) = \prod_j r_j$ by a gate h with (i) h twice continuously differentiable on $[0, 1]^{n+1}$, non-decreasing in each r_k and with $\partial^2 h / \partial r_i \partial r_j > 0$ for $i \neq j$ on the interior, (ii) $c_1 \prod_j r_j \leq h(r) \leq c_2 \prod_j r_j$ for constants $0 < c_1 \leq c_2$, and (iii) $h(1, \dots, 1) = 1$. Then T1 to T5 and Propositions 1 to 3 hold verbatim, with

a_1^{n+1} replaced by $c_1 a_1^{n+1}$ in T3(c2).

Proof. The gate enters at five places. Invariance of S and Proposition 2 use only $\chi_{\pm}(0) = 0$ together with the fact that the gate takes both zero and positive values on S , which (ii) supplies because the zero sets coincide. T4 is unchanged with $N(t) = \int_0^t \chi_+ h \, ds$. The complementarity clause of T1 is exactly (i), which is why the differentiable form is required there rather than finite-difference supermodularity; the reference gate of the checks, $h = (\prod_j r_j)(1 + r_3)/2$, satisfies it with $c_1 = 1/2$ and $c_2 = 1$. The persistence estimate of T5 needs $\liminf h > 0$, which follows from the lower bound in (ii) and the persistence bound already proved for the product; the collapse side needs $h \rightarrow 0$, which follows from the upper bound. T2 needs $h(1, \dots, 1) = 1$ and $h = 0$ at the failed corner, which are (iii) and (ii). The construction is therefore not tied to the product form but to the zero set and to two-sided comparability, which is also the sense in which Lemma C1 identifies the slow channel count with the gate.

A leak changes one statement, not the theorems. With $h_{\delta} = \delta + (1 - \delta)\prod_j r_j$ and $\delta > 0$ the zero sets differ. S is still invariant, since $\chi_{\pm}(0) = 0$ does not involve the gate. For selectors in the exclusive class E , both corners remain equilibria and the dichotomy is unchanged, because at the failed corner the accumulation channel is switched off by χ_+ rather than by the gate. This last statement is specific to the exclusive class. At finite

sharpness $\chi_+(-a_0) > 0$, so a leak makes $dQ/dt = \varepsilon\chi_+(-a_0)\delta > 0$ at the failed corner and

E_{path} is displaced as well: on the invariant face $r = 0$ the equilibrium solves $\delta\chi_+(-a(Q))(1 - Q) = \rho\chi_-(-a(Q))Q$, that is $Q/(1 - Q) = (\delta/\rho)\exp[-\gamma a(Q)]$, so $Q \leq (\delta/\rho)\exp(-\gamma a_1)$.

Allowing a leak and finite sharpness together therefore displaces both corners by explicitly small amounts, and is a different perturbation from Proposition 6. What fails is the exact halting of accumulation: while a prerequisite is offline the evidence rate is $\varepsilon I_0 \delta$ rather than zero, so the statement that a single offline prerequisite halts accumulation weakens to reducing it to the leak rate. At $\delta = 0.01$ the switching boundary is unmoved to within $2e-8$ and the two outcomes are unchanged, while dQ/dt at a configuration with one capacity offline is $2.8e-7$ in place of exactly zero.

Boundary drift when E4 is violated. Replace λ by $\lambda_\eta = \eta + (1 - \eta)L(w_+)$, so that $\chi_\pm(0) = \eta/2 > 0$ and the hypothesis of Proposition 2 fails by η . On S the normal velocity is $dz/dt = -a'(Q)dQ/dt = -a'(Q)\varepsilon(\eta/2)[C(r)(1 - Q) - \rho Q]$, which is $O(\varepsilon\eta)$, so the switching boundary moves by $O(\varepsilon\eta)$ to first order. Numerically, at $\Delta = 0.55$ the boundary exceeds Q_c by $9.67e-4$ and $4.68e-3$ for $\eta = 0.05$ and 0.25 at $\varepsilon = 0.02$, and by $2.43e-4$ and $1.21e-3$ at $\varepsilon = 0.005$; the ratio of the drift to $\varepsilon\eta$ lies between 0.94 and 0.97 throughout. The exact invariance of S is therefore the $\eta = 0$ member of a continuous family and not a fine-tuned point: a system that suspends updating only approximately at ambiguity has a

switching surface displaced in proportion to $\varepsilon\eta$, which is small whenever the evidence is slow.

Verification

Positive invariance was checked on the boundary of the state cube, where the field points inward³⁸. The comparison inequality used by Proposition 6(b), $g(x, Q) \leq K^*x$ on $[0, 1]$, was checked on a grid, the maximum of $g(x, Q)/x$ being 0.1806 against $K^* = 0.2908$. Both equilibria were confirmed exact ($dQ = 0$), and the finite-difference spectra at the two corners, taken one-sided on the faces of the state box so that no arm of the difference is pinned by the clip to $[0, 1]$, agree with the analytic Jacobian to within $1e-9$: the state diagonals are $-(1 - a_1) - \kappa_k$ and $-a_0 - \kappa_k$, and the evidence directions are $-\varepsilon\chi_+(1 - a_1) = -0.005011$ and $-\varepsilon\rho\chi_-(-a_0) = -0.001244$, the latter vanishing at $\rho = 0$ and negative for $\rho > 0$. Invariance of S was confirmed by the collapse boundary from a recovered start equalling Q_c at $\varepsilon = 0.02$ and $\varepsilon = 0.005$ alike. Propagation was verified well below the sufficient bound and to fail at $\kappa = 0$; absorption was verified from the recovered configuration and shown to fail for a root-recovered, downstream-pathological start at low coupling. The evidence law of T4 was checked against the augmented system carrying N as a coordinate, and the integrating-factor solution against a finite- γ trajectory on which both channels act, in both cases to within $1e-12$.

Proposition 1, which is proved for the whole E4 class in the Methods, was corroborated numerically on three members: the class conditions $L(1/2) = 0$ and $L(0) = L(1) = 1$ were verified, $\chi_{\pm}(0) = 0$ at every γ with the collapse boundary at Q_c for every ε , and $\sup|\lambda_{\gamma} - 1|$ on $\{|z| \geq 0.05\}$ was seen to fall to below $1e-3$ by $\gamma = 500$ for each member, which is the uniform convergence used by part (b). The unit jump of the limit pair at $z = 0$, which is what makes the regularisation necessary, was confirmed, and with E4 dropped the boundary drifts and returns towards Q_c as ε decreases. T5 was checked by convergence of off-manifold trajectories to the corner selected by the sign of $z(0)$ at $\kappa = 0.6$, endpoint data included, by convergence of orbits started on S to the corresponding point of E_s , and by stable mixed equilibria for $\kappa < a_1^2/4$ with the saddle-node at $\kappa_1 = a_1^2/4$.

Proposition 3 was checked by ordered first-passage times and by the synchronous recovery of a uniform onset. Proposition 6 was checked at $\gamma = 10, 20$ and 40 from a recovered start: the state converges to $(1,1,1,1)$ in each case, the explicit margin of Lemma F at $\gamma = 20$ is $\bar{z} = 2.49e-2$ while the observed margin never falls below 0.4125 , and $1 - Q$ agrees with the fixed point $pe^{-\gamma(1-a(Q))}Q$ to a relative $1e-3$, taking the values $2.04e-4, 4.14e-8$ and $2.33e-15$ against the bounds $1.83e-2, 3.35e-4$ and $1.13e-7$; the collapse side reaches E_{path} to machine zero. Proposition 7 was checked on the comparable gate $h = (\prod r_j)(1 + r_3)/2$, which leaves the boundary at Q_c to $2e-8$ and both

outcomes unchanged, and on the leaky gate at $\delta = 0.01$, which does the same while giving $dQ/dt = 2.8e-7$ with one capacity offline. The drift under a violation of E4 was measured on the four combinations of ε in $\{0.02, 0.005\}$ and η in $\{0.05, 0.25\}$, the ratio of drift to $\varepsilon\eta$ staying between 0.935 and 0.973. These switching-boundary measurements are bisected to $1e-7$ in every run mode, the drifts being of order $1e-4$ to $1e-3$, so the quoted values do not depend on the mode; the coarser sweeps selected by the fast mode apply to the other checks only. A leak at finite sharpness was checked separately: at $\gamma = 20$ and $\delta = 0.01$ the failed corner moves to $Q = 6.14e-8$, matching the predicted solution of $Q/(1 - Q) = (\delta/\rho)\exp[-\gamma a(Q)]$ to twelve digits and lying below the bound $(\delta/\rho)\exp(-\gamma a_1) = 4.98e-4$. Proposition 5 was checked on the worked instance: the informative regime holds, $\partial G_1/\partial r$ and $\partial G_1/\partial \Lambda$ are negative over the sampled range with maxima $-9.3e-3$ and $-3.1e-6$, the derived threshold runs from $a_0 = 0.60$ at $Q = 0$ to $a_1 = 0.15$ in the limit, its first and second derivatives in Q are everywhere negative at $\Lambda^* = 2$, the analytic elasticity falls from $q(\Lambda_0) = 3.91$ to 2, and the criterion $\max q(\Lambda)\Lambda^*/\Lambda$ takes the values 0.78 at $\Lambda^* = 2$ and 1.95 at $\Lambda^* = 5$, the second case being convex over part of the range, which is the equivalence of part (b). Proposition 4 was checked by implementing the companion observation model as five Poisson channels: the joint Fisher information about πv_{slow} agreed with the slow-channel mean to a relative $2.5e-10$,

the other four channels showed exactly zero dependence on that parameter, the ratio $M_{\text{slow}}/C(r)$ stayed within $[1.01, 1.94]$ against the bound $K = 2.31$ of Lemma C1, the expected count of usable per-context estimates matched v times the integral of $C(r)$ to a relative $7e-6$ over $2e5$ sampled contexts, and at $\varepsilon = 0.001$ the evidence coordinate reached the common threshold at $t = 1.6e3$ against 3.7 to 7.5 for the four capacities, which is the corollary. Table 1 was illustrated one row at a time.


Scope and premises

It is not claimed that every adaptive system has this form, that dependency structure by itself generates bistability, that the structure can be recovered from trajectories, that the coupling bound is sharp, or that the theorems cover branching graphs. Cyclic structures and backward coupling lie outside the class, the forward ordering being used wherever the cascade is triangular.

Tables 1 and 2 list the status of each result. The observation-model premise, the bridging conditions O1 to O3 used by Proposition 4, the decision-layer premises D1 to D4 used by Proposition 5, the premises E1, E2 and E4, the sharpness of the decision and the regularisation E are modelling commitments; the theorems are conditional on them, and the Discussion states what each buys and how it can fail. T5 is proved only under the explicit strong-coupling condition; outside it the classification of all attractors

remains open. The open empirical items are prospective tests of the observation model, of the sharp versus graded decision, of the coupling threshold and of the consolidation time.

Supplementary Figure 1



Supplementary Figure 1. Numerical corroboration of the four claims that

distinguish the construction. a, The accumulation selector χ_+ . The identifiability

weighting of E4 forces $\chi_+(0) = 0$ at every sharpness γ , whereas plain model averaging

leaves $\chi_+(0) = 1/2$; the reference regularisation $\exp(-s/z^2)$ with $s = 0.01$ is shown for

comparison. **b**, A recovered configuration receiving an excursion to $r_0 = 0.45$ on either side of the switching surface, at $Q_c \pm 0.002$: the root and the evidence separate to the two corners, which is the dichotomy of T5. **c**, The collapse boundary from a recovered start against the evidence timescale. With E4 it sits at $Q_c = 0.5774$ for every ε , so S is invariant at finite sharpness; without it the boundary drifts and returns to Q_c only as $\varepsilon \rightarrow 0$. **d**, The low first-stage root at $Q = 1$ as a function of κ_1 . The two roots coalesce at the saddle-node $\kappa_1 = a_1^2/4 = 0.005625$, where the low root is no longer hyperbolic.