Supplementary Information

The Birth, Life, and Death of Internal Organization

Yuxuan Deng ORCID: 0009-0000-1531-2152

Version 2

Scope and statistical units

The supplementary material expands already executed Paper I evidence. Tables S1–S4 retain the original observations and negative results. Note S5 preserves the exploratory dynamics and its quantitative predictive boundary. The separate three-page V19 evidence document provides the full-cohort visualization. No new core experiment was run for this revision. Source tables accompany the manuscript; repeated branch or checkpoint records are not additional histories.

Table S1. Fresh architecture confirmation

Architecture	Seed	Form	Reorg	Death	Persist	Full	Null/B	Continued/B
mlp	4000	yes	yes	yes	yes	yes	0.012771	0.97466
mlp	4001	yes	yes	yes	yes	yes	0.00089717	1.3449
cnn	4010	yes	yes	yes	yes	yes	0.0025814	0.59831
cnn	4011	yes	yes	yes	yes	yes	0.00048261	1.2107
gru	4020	yes	no	yes	yes	no	0.044403	0.77937
gru	4021	yes	yes	yes	yes	yes	0.043923	0.8652
transformer	4030	yes	yes	yes	yes	yes	0.047059	1.2254
transformer	4031	yes	yes	yes	yes	yes	0.024908	0.94107
gnn	4040	yes	yes	yes	yes	yes	0.01245	1.2566
gnn	4041	yes	yes	yes	yes	yes	0.0022667	0.93547

Each row is one history condition. Null and continued branches share its B endpoint. The failure at GRU 4020 remains visible. No-migration findings refer only to the two sampled candidate sites.

Table S2. Paired original and transformed coordinate tests

Kind	Architecture	Form	Transformed	Full	Corr. A	Error A	Error B
stable	mlp	yes	yes	yes	1	2.3246e-06	2.3842e-06
stable	cnn	yes	yes	yes	1	1.9372e-06	2.3842e-06
stable	gru	yes	yes	yes	0.99977	5.6028e-06	2.7418e-06
stable	transformer	yes	yes	no	0.7816	2.0862e-06	2.265e-06
stable	gnn	yes	yes	no	1	1.4305e-06	2.5034e-06
control	mlp	no	–	–	–	–	–
control	cnn	no	–	–	–	–	–
control	gru	no	–	–	–	–	–
control	transformer	no	–	–	–	–	–
control	gnn	no	–	–	–	–	–

Original and transformed readings are paired within model. The five shuffled rows are matched controls; – denotes an unmeasured field, not a failed test. Functional errors use the same sampled-input test, not a proof over all inputs.

Table S3. All eighteen fresh V19 history conditions

Seed	Eligible	Null role/B	Null energy/B	Continued role/B	Continued energy/B	Death
3100	yes	0.0011501	0.28963	1.0347	1.02	yes
3101	yes	1.1473e-16	9.0342e-15	0.66551	1.1912	yes
3102	yes	0.0066474	0.64167	0.97128	0.89764	yes
3103	yes	0.00085486	0.37899	0.93892	1.4636	yes
3104	yes	7.493e-06	0.00052474	0.76611	1.3246	yes
3105	no	0.010497	0.26	0.78806	0.75514	not eligible
3106	no	0.013631	0.78687	1.2731	1.075	not eligible
3107	yes	0.010584	1.8923	0.89781	0.8105	yes
3108	no	0.039161	0.6015	0.93252	0.91066	not eligible
3109	yes	0.0062146	1.126	1.0333	1.1856	yes
3110	yes	6.0859e-07	7.6742e-06	0.96763	1.0159	yes
3111	yes	0.028841	0.97359	1.1271	1.0498	yes
3112	yes	0.031521	0.3942	1.0758	0.94154	yes
3113	yes	4.2415e-06	0.00026053	0.94733	1.0521	yes
3114	yes	0.010214	1.4068	0.93704	1.0326	yes
3115	no	0.40209	0.73221	0.86276	0.87898	not eligible
3116	no	0.090066	1.4218	1.0647	0.62131	not eligible
3117	no	0.040314	0.10799	9.0593	0.84475	not eligible

Each row is one source history condition, with paired terminal branches at step 3200. Six source-ineligible histories retain measured values but do not count as eligible death tests. Carrier energy is activity in the frozen carrier coordinates, not an independent memory-content certificate.

Table S4. Continuous forecast compared with the archived baseline

Branch	Normal-form MAE	Early-hold MAE
Null	0.0189744	0.0170958
Continued	0.333066	0.203940

Both estimates use the same twenty-two selected histories. The continuous-error comparison favors the early-hold baseline. Correct ordering or ratio-classification does not remove this quantitative limitation.

Supplementary Note S5. Exploratory lifecycle dynamics and its predictive boundary

The lifecycle gates above were defined and frozen independently of any dynamical model. A paired-update decomposition provides exploratory bookkeeping for the observed transitions, not an independently established neural predictive law. At a fixed audit resolution and within one realization, let $\mathfrak{C}_t$ denote the intervention–consequence map after update t, and let Δ_t denote the currently supported consequence pressure. Relative to a paired evolution in which that pressure is absent, write

$$\boxed{\mathfrak{C}_{t+1} - \mathfrak{C}_t = R_t(\Delta_t) + U_t(\mathfrak{C}_t)}. \tag{S1}$$

Here R_t is the consequence-dependent recruitment or redirection term—the change attributable to supported consequence relative to its paired pressure-removed evolution—and U_t is the evolution of the same organizational realizer in the absence of that pressure. The decomposition is exact bookkeeping for the paired update; it is not, by itself, a lifecycle theorem. Predictive content enters through independently testable constitutive conditions on recruitment, contraction and leakage, together with their finite-time consequences.

The dynamics acts on $\mathfrak{C}_t$, not directly on $\mathscr{O}_Q$. After each update the local quotient $E_t^{(r)}$ is induced from the updated consequence map, and consequence-preserving identification across compatible realizations yields $\mathscr{O}_Q(t)$. A recruitment term that separates two previously consequence-equivalent interventions produces a fiber split and therefore birth. Continued support can preserve the split. Changed consequence can redirect the anchored incidence or reverse its sign. When support is removed and unsupported evolution contracts the active consequence contrast, formerly distinct interventions merge under the frozen audit and the organization dies. Leakage or overlapping support can instead produce coexistence or a nonzero residual floor.

The constitutive conditions are non-vacuous. Supported consequence without sufficient recruitment need not produce formation; removing support need not produce death when leakage or persistence dominates; and a one-step contraction estimate does not by itself certify contraction throughout an entire release phase. These excluded regimes are observed in the control record and prevent the dynamical form from merely restating the lifecycle labels after the fact. In transparent adaptive classes, constitutive quantities can be estimated before an unseen phase and used in frozen prediction of lifecycle regime, but quantitative transfer requires realization-specific calibration. The present deep-neural and naturalistic experiments do not claim parameter-free universal kinetics or prospective recovery of all phase-domain coefficients from a compact state description.

What the archived early-trajectory fits predict

An archived normal-form analysis selected 22 histories that had already formed and reorganized from the preceding V20, fresh V20b and whitened V21 studies (9, 9 and 4 histories, respectively). It therefore reanalyses overlapping cohorts rather than supplying 22 new confirmations. Fitting $r(t) = r_\infty + (1 - r_\infty)e^{-kt}$ through step 200 gave the correct terminal Null-versus-continued ordering in 22/22 pairs and the correct ratio-only Null death classification in 21/22. The 21/22 figure counts correct classifications, not observed deaths: two histories correctly retained a non-death classification. These ratio-only predicates are weaker than the full lifecycle gate.

Quantitative prediction was less successful. Saved endpoint mean absolute errors were 0.0189744 for Null and 0.333066 for continuation, compared with 0.0170958 and 0.203940 for simply retaining the early observed value. The fit therefore did not improve continuous endpoint prediction over that baseline. It supplies an explicit qualitative comparison, alongside a quantitative limitation. A separate later frozen cross-mechanism forecast also failed on its tested family; it is retained as a negative methodological boundary and does not supply a universal constitutive law.

Source clarification: HAPT decision threshold

The hash-locked executable configuration specifies a death fraction of 0.25, consistent with the confirmation report and saved seed-level decisions. An earlier overview document stated 0.30. The configuration and audit code match their recorded protocol-lock hashes; the 0.25 decision and original 2/5 full lifecycle result are retained. The discrepant overview is not used to change the decision after observing outcomes.