The Boundary That Holds

A Falsifiable Model of Organismic Integration

Madhav Moganti
Bell Labs, Nokia — mmoganti@outlook.com

Abstract

Living organisms stay coherent wholes even though their molecules turn over, their control is distributed across many systems, and they constantly exchange matter and energy with their environment. How they manage this is usually described only in metaphor — as a "self," an "organizing principle," or a "hidden controller." This paper proposes a way to study it without assuming any of that.

Organismic Integrative Completeness (OIC) breaks organismal wholeness into five measurable components — how well an organism's subsystems participate in a shared network, how completely a compact organism-level state can represent them, how far its causal influence reaches, how long its coordination persists, and how reliably the organism recovers after disruption — each with an explicit, computable estimator rather than a metaphor. A stronger hypothesis, the Boundary–Membrane Integrative Core (BMIC), proposes that a single low-dimensional variable performs selection, compression, and coordination across an organism's genomic, bioelectric, endocrine, immune, and neural systems, and states one concrete statistical law linking that variable to measurable subsystem activity.

The paper builds in ways for BMIC to fail: a cross-fitting procedure that keeps the model from "predicting" a subsystem using information secretly drawn from that same subsystem, standard statistical controls, and a fully reproducible synthetic example with published code — extended to a preregistered comparison against a more demanding nonlinear competitor, since a weak baseline would make almost any model look good. A four-tier claim structure keeps modest, already-defensible claims separate from the far stronger, currently untested claim that this coordinating variable is anything more than a useful summary of ordinary physical processes; this paper does not make that stronger claim. A companion paper extends the same framework, under a stricter evidential bar, to the further question of self-reference and subjective experience.

Keywords: organismic integration; network physiology; causal emergence; falsifiability; latent-state modeling; boundary participation; biological complexity; developmental bioelectricity

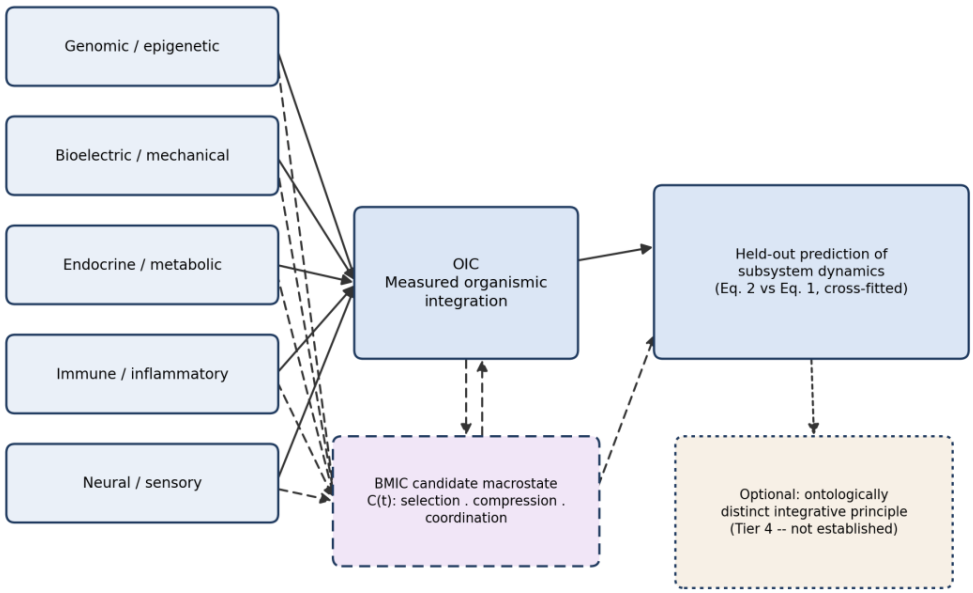

Figure 1. Established physical systems feed measurable OIC constructs (solid arrows). BMIC proposes a candidate macrostate C(t) with hypothesized, not-yet-established coupling to those systems (dashed). Whether that macrostate is more than a coarse-grained summary variable is a further, unestablished question (dotted) that this paper does not address.

1. Introduction

A living organism is not a fixed parcel of matter. Its molecules turn over, its cells divide and die, its organs shift state, and the information it receives from moment to moment never repeats. Yet the organism holds a boundary, coordinates its parts, recovers from injury, and acts as one continuing thing. Classical work on morphogenesis, self-organization, autopoiesis, and homeostasis showed that this kind of order can emerge from distributed physical processes, with no need for a central assembler [1–4]. More recent work on network physiology shows that the way organs interact with each other shifts with physiological state, and that a breakdown in that coordination can cascade through the whole organism [5,6].

None of this establishes a controller of any kind. What it establishes is distributed organization. This paper asks three narrower questions instead. Can organism-level integration be turned into something measurable, rather than left as a metaphor? Can a single candidate variable be defined precisely enough to make falsifiable predictions — including the prediction that it doesn't exist? And can that variable be specified so concretely that a worked synthetic example, and eventually real data, could confirm or rule it out?

The answer developed here has two layers. The first, Organismic Integrative Completeness (OIC), makes no ontological commitment at all: it simply asks whether an organism's subsystems participate in a shared coordinated domain, whether a compact organism-level description captures what's needed for viable action, whether that description's influence reaches the right effectors, whether the coordination persists over time, and whether the organism can recover from disruption. The second, the stronger Boundary–Membrane Integrative Core (BMIC) hypothesis, proposes that a single low-dimensional variable, C(t), actively performs the selection, compression, and coordination behind that first layer, linking an organism's genomic, bioelectric, endocrine, immune, and neural systems. BMIC can be read two ways: modestly, as a useful summary variable for prediction and control, or ambitiously, as evidence of something ontologically

distinct. Only the modest reading is developed in this paper — evidence for it would not, by itself, support the ambitious one.

That restriction is a choice about evidence, not a way of avoiding the more interesting question. Whether this kind of coordination has anything to do with self-reference or subjective experience in organisms with nervous systems is a real question, and a companion paper takes it up directly, under a stricter evidential standard than anything argued here [25]. Every claim in this paper can be evaluated, and could fail, without that further question ever coming up.

2. Claim Structure and Evidential Boundaries

To keep these different kinds of claims from bleeding into each other, the paper sorts them into four tiers (Table 1). Higher tiers require more evidence, but they are cumulative only in evidential burden, not in truth — a reader can accept Tier 1 mechanisms while rejecting every tier above it.

Tier	Claim type	Status and admissible inference
1	Measured physical mechanisms	Genomic regulation, cell signaling, tissue mechanics, bioelectric patterning, endocrine and immune signaling, and neural dynamics.
2	Organism-level construct	A coarse-grained state (like OIC) is useful if it improves prediction, control, or classification across subsystems. Assumed physically realizable unless evidence says otherwise.
3	BMIC coupling	A preregistered global-state variable adds predictive power beyond the best physical null model (linear or nonlinear, §6), on held-out data, and survives intervention and replication.
4	Ontologically distinct principle	The candidate state is more than a summary of physical dynamics — evidence beyond prediction alone would be required. Not addressed by this paper.

Table 1. Claim classes used throughout the paper.

This paper develops Tiers 2 and 3 in full and does not attempt Tier 4. The companion paper takes up a narrower version of Tier 4 — whether self-referential integration bears on subjective experience — under its own evidential standard; the two papers are independent, and neither one's results are required to evaluate the other [25].

3. Empirical Constraints on an Organismic Integrator

3.1 Coordination begins before there is a nervous system to do it

A human embryo starts as a single organized cell carrying parental genomes and maternally supplied cytoplasm, RNA, proteins, and organelles. Well before implantation, embryos — including stem-cell-based embryo models — already show substantial self-organization [13,14]. The genome activates around the eight-cell stage [15], and everything that follows depends on transcriptional state, chromatin accessibility, position, signaling history, and mechanical forces [16–19]. Developmental bioelectricity adds another experimentally manipulable channel here: membrane voltage and gap-junction coupling between cells can change anatomical outcomes, and bioelectric information is integrated across developing tissue well before neurons exist [9–12].

This sets a hard constraint on BMIC. If it contributes anything to how an organism is built, its developmental interface cannot run exclusively through a mature nervous system — a pre-neural candidate has to be

specified in cellular, bioelectric, mechanical, or molecular terms that can actually be measured. And because these mechanisms already explain most of morphogenesis on their own, BMIC only earns its place if a preregistered residual survives increasingly complete physical models and direct intervention.

3.2 Integration is distributed, and it shifts with physiological state

Homeostasis is active regulation within a viable range, not a static equilibrium [4]. Network-physiology research shows that the relationships among brain, heart, lungs, and other organ systems reorganize as physiological state changes [5,6] — which supports an organism-level description but doesn't require a single central store of information. Constraint-based, homeostatic, active-inference, and causal-emergence accounts are BMIC's most serious competitors precisely because they already explain organismal wholeness through distributed physical organization alone [7,8,20–23]. Throughout this paper, BMIC is treated as a hypothesis that has to outperform these accounts, not one that replaces them.

3.3 Losing integration is a limit of the model, not evidence of something else

OIC's recovery component, c^V, makes a specific, testable prediction at the edge of life: as coordination among subsystems fails, the ability to recover from disruption should degrade gradually and measurably, not drop off a cliff at some arbitrary threshold. At irreversible death, most of the organism's matter and anatomy is still present, but energy production, membrane gradients, circulation, repair, and multiorgan regulation all fail together. That transition is evidence that recoverability and integration have been lost. It is not evidence that some separate entity has departed. Operationally, death is simply the failure to return to the viability region that c^V measures — and any stronger claim, that some distinct integrative substrate persists or transforms rather than just ceasing to be measurable, would be a Tier 4 claim this paper does not make.

4. The Boundary–Membrane Integrative Core Hypothesis

4.1 What BMIC claims, precisely

BMIC proposes that every living organism instantiates an organism-level integrative state, built from two functions. A boundary operator selects, aggregates, and compresses the information relevant to the organism as a whole. A representational membrane maintains that information as a single low-dimensional state, C(t), linking the organism's physical subsystems, its recent history, and the actions available to it. Local, specialized systems still hold their own information and run their own operations — BMIC predicts two-way coordination between them and the candidate global state, not that they hand their information over to it. A stronger reading would treat C(t) as ontologically distinct from the physical subsystems it coordinates. That reading isn't assumed here, and nothing in the operational model below depends on it.

4.2 Making "boundary" and "membrane" concrete

"Boundary" here doesn't mean an anatomical enclosure or isolation from the environment. It means an organism-relative partition defined by which subsystems selectively participate in, and reciprocally influence, a shared causal network. Table 2 turns this into a specific, computable quantity: build a causal network among the organism's measured subsystems, and ask what fraction of them fall inside its largest strongly connected component — the largest set of nodes that can all reach one another by some directed path. That fraction, benchmarked against randomized and fully connected reference networks, is boundary participation (c^B), the estimator used throughout the rest of this paper. "Selective participation and reciprocal influence" isn't a placeholder for a later definition; strongly-connected-component membership is that definition.

"Membrane" names the representational function that links subsystem state, environment, history, and available action inside the candidate macrostate. Boundary and membrane are functional roles, not anatomical organs waiting to be discovered — their value depends entirely on whether each role can be operationalized and tested against simpler alternatives.

4.3 Where this fits among known systems

The genome, epigenome, bioelectric and mechanical networks, endocrine and metabolic systems, immune system, and (where present) nervous system are all established physical information systems in their own right. BMIC doesn't claim to relocate their information into one central store; it proposes a coarse-grained state that captures organism-level variables none of them holds alone. The real test, then, is whether that macrostate adds reproducible predictive or causal power beyond what distributed physical coordination already provides.

The developmental interface has to change with age. Before a nervous system matures, any candidate global state has to be expressed through cellular, bioelectric, mechanical, and chemical organization alone. Once a nervous system exists, the highest-bandwidth interface is more likely distributed across neural, endocrine, and autonomic loops than concentrated in one structure — a given subsystem might be necessary for coordination without being sufficient for it on its own.

4.4 A coupling law precise enough to fail

Here is the same idea in a form specific enough to test. Let B(t) be a snapshot of the organism's measured physical subsystems at time t (say, one summary number per subsystem in Figure 1), E(t) the environmental inputs, and ε(t) ordinary process noise. The skeptical baseline — no coordinating variable at all — is a standard linear model of how each subsystem evolves from the last:

$$H0:\ B(t+1) = A \cdot B(t) + G \cdot E(t) + \varepsilon(t) \quad (1)$$

BMIC's alternative adds one thing to this baseline: a single low-dimensional variable C(t), entering through a coupling strength μ. The simplest version of that coupling is linear, μ·C(t):

$$H1:\ B(t+1) = A \cdot B(t) + G \cdot E(t) + \mu \cdot C(t) + \varepsilon(t) \quad (2)$$

$$C(t+1) = \rho \cdot C(t) + \Gamma \cdot S(t) + \eta(t), \quad \text{where } |\rho| < 1 \quad (3)$$

In words: each subsystem's next state depends on its own recent history (A·B(t)), the environment (G·E(t)), and — under H1 only — a pull from the shared variable C(t). That shared variable updates itself from preregistered source information S(t) drawn from the physical subsystems, with ρ setting how long its influence lingers after a subsystem stops driving it. (A bounded alternative, μ·tanh[C(t)], is available if saturating outputs are needed.) Because μ is a small, bounded, estimable quantity — not an unconstrained free function — it can be tested against zero, and any model comparison has to penalize the extra parameters using held-out prediction error. If C(t) doesn't improve prediction or the organism's response to intervention, the coupling hypothesis loses. Equations 1–3 are the only definition of C(t) used anywhere in this paper; the worked example in §6.2 and the code in Appendix B implement this exact system, term for term, with nothing added or dropped.

4.5 How to perturb it

Any manipulation that reliably shifts subsystem coordination can test the roles assigned to the membrane. Good candidates — anesthetic or pharmacological state changes, immune challenge, metabolic or nutritional stress, developmental intervention — share three properties: they're repeatable, already have a measurement literature, and directly target one or more of the subsystems in Figure 1. The prediction isn't that a perturbation will unmask some hidden membrane. It's more modest: a well-specified manipulation might shift one or more measured subsystem states, change the temporal weighting captured by ρ and Γ, or

improve C(t)'s stability and held-out predictive accuracy — all while staying fully compatible with ordinary physiological plasticity. A genuinely BMIC-specific result needs one more step: showing that an independently estimated C(t) improves held-out prediction or intervention response beyond what physical, physiological, and dose-dependent models already explain.

4.6 Avoiding an infinite regress of readers

Any representational state invites a specific objection, one Dennett raised against "Cartesian materialism" [24]: if you name a place or state where information "comes together," you risk smuggling in an implicit further reader — some additional system the assembled representation is displayed to. That just relocates the explanatory burden to how that further system uses the display, rather than discharging it. This is a general risk for any account of coordination, with nothing to do with experience specifically: if C(t) has to be read by some other component before it does any causal work, the model hasn't explained coordination, only postponed explaining it.

This paper avoids that trap by construction. C(t), as defined in Equations 2–3, is the coordinating event itself, not a display waiting to be interpreted — it enters directly into the dynamics of B(t+1) in Equation 2. Nothing else has to read it first. Were that not true, the model would have merely relocated the coordination problem, and Section 4.2's claim that boundary and membrane are functional roles rather than hidden organs would collapse — an implicit "reader" organ would be doing exactly the work being denied. None of this bears on whether occupying the state in Equations 2–3 is also sufficient for experience in organisms that have it; that further, narrower question belongs to the companion paper [25], not to this one.

5. Organismic Integrative Completeness

5.1 Definition

OIC is functional, organism-relative, and deliberately modest in what it claims. It doesn't mean omniscience, storing every microscopic variable, logical completeness, or exemption from ordinary physical limits. An organism is integratively complete to the degree that its relevant subsystems participate in a coordinated domain, its global state captures what's needed for viable action, that influence reaches the right effectors, the coordination persists over time, and the organism can recover after a standardized disruption.

Table 2 turns each of these five ideas into a number. Representation (c^R), for instance, is the organism-level counterpart of Equations 2–3: it's high to the extent that the organism-level state can reconstruct subsystem-level variables without peeking at the outcomes it's meant to predict. Like the other four components, it's a single aggregate score — useful for comparison, but not a replacement for the underlying subsystem measurements it's built from, and uninformative if those measurements haven't themselves been validated.

Component	Operational definition	Candidate estimator, normalized to [0,1]
Boundary participation, c^B	Extent to which an organism's subsystems participate in one reciprocally connected causal network.	Weighted share of relevant nodes in the largest strongly connected component, benchmarked against randomized and fully connected reference networks.
Representation, c^R	Ability of the candidate global state to reconstruct a preregistered organism-level variable without using the target outcome.	Held-out R^2, floored at zero and scaled to the achievable reliability ceiling.

Component	Operational definition	Candidate estimator, normalized to [0,1]
Causal reach, c^c	Breadth and specificity of subsystem responses, in the predicted direction, to intervention on the global state or a valid instrument.	Weighted share of target systems with preregistered causal effects in the predicted direction, penalized for indirect-path violations.
Temporal continuity, c^T	Persistence and history-dependence of the organism-level state across a specified interval.	Normalized predictive information between the state now and later, estimated on held-out time blocks.
Viability recovery, c^V	Probability of returning to a preregistered viability region after a standardized disruption.	Empirical recovery probability within a fixed time window, adjusted for disruption severity.

Table 2. Operational definitions and normalization protocol for the five OIC components.

Every component has to be defined in advance, estimated on held-out data, and reported with uncertainty intervals. Combining them is deliberately conjunctive rather than additive:

$$OIC(t) = \prod_j c_j(t)^{w_j}, \quad j \in \{B, R, C, T, V\}, \quad \Sigma_j w_j = 1 \qquad (4)$$

A geometric mean means a high score on one component can't paper over failure on another indispensable one — equal weighting is a reasonable default for an initial study, but domain-specific weights need their own validation, and an arithmetic mean and a minimum-component rule are worth reporting alongside it as sensitivity checks.

5.2 A worked example

Figure 2 makes the arithmetic concrete with a synthetic, hand-assigned example — not data from a real organism, just a demonstration of how the index behaves (full values in Appendix A.1). Give an "integrated" state high scores across all five components, and OIC comes out to 0.839. Degrade causal reach, continuity, and recovery — leaving boundary participation and representation relatively intact — and OIC for that "perturbed" state falls to 0.504. A threshold like 0.65 is illustrative only; using one for real classification would require prospective calibration.

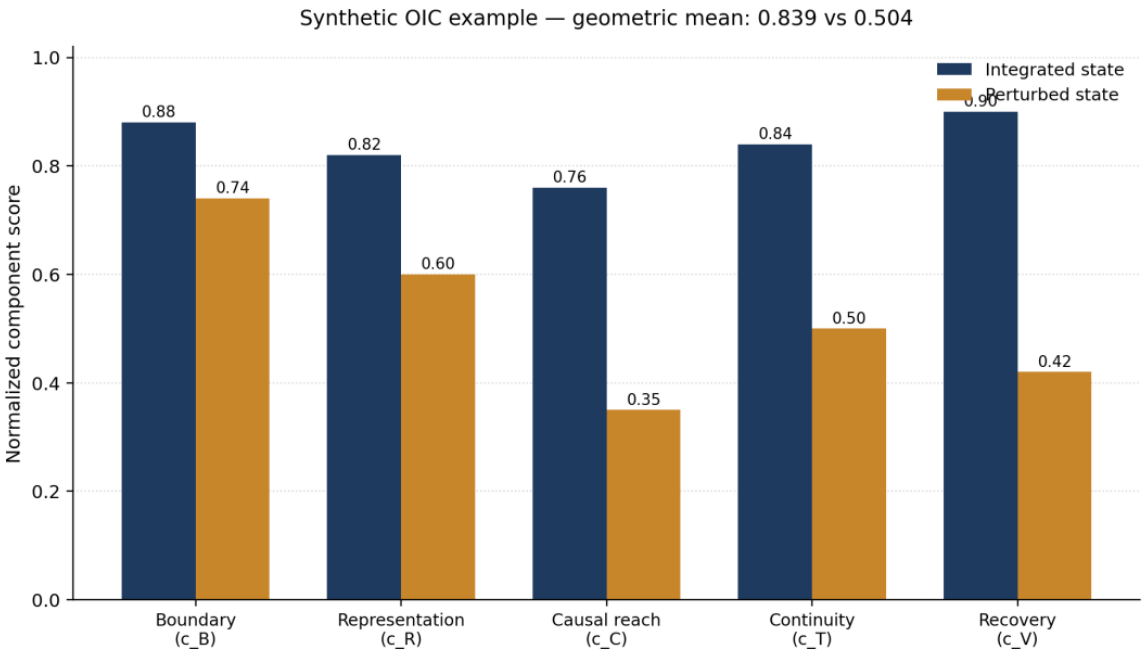

Figure 2. The geometric mean is conjunctive: jointly degrading causal reach, continuity, and recovery pulls the perturbed state's score down more than any single weak component would on its own. Values are illustrative only, not biological or clinical estimates.

6. Noncircular Inference Design and a Nonlinear Competitor

6.1 Ruling out circularity

A latent state built from the same target it's used to predict will look informative no matter what — that's a construction artifact, not evidence. To rule it out, C(t) is estimated with leave-one-system-out cross-fitting: for a given target subsystem T, the encoder that builds C(−T,t) is trained only on source data that excludes T and anything directly downstream of it, then frozen before evaluation. Prediction is tested on held-out subjects or non-overlapping time blocks, with the target role rotated across every subsystem in turn. Time-shifted, phase-randomized, and cross-subject permutation controls (Figure 3, Stage A) estimate how much apparent improvement comes from ordinary autocorrelation and model flexibility alone, so it can be subtracted out.

Even a design this careful can only establish a useful macrostate — it can't show that the macrostate is anything more than a coarse-grained physical summary. A stronger coupling claim needs an independent measurement Y(t), not built from the target network, or a randomized instrument U(t) that shifts the candidate interface while satisfying a defensible exclusion restriction (Figure 3, Stage B) — and any direct physical path from U to the target has to be measured and modeled explicitly. Without that independent channel, residual predictive power is evidence of some omitted structure, not evidence of a distinct macrostate.

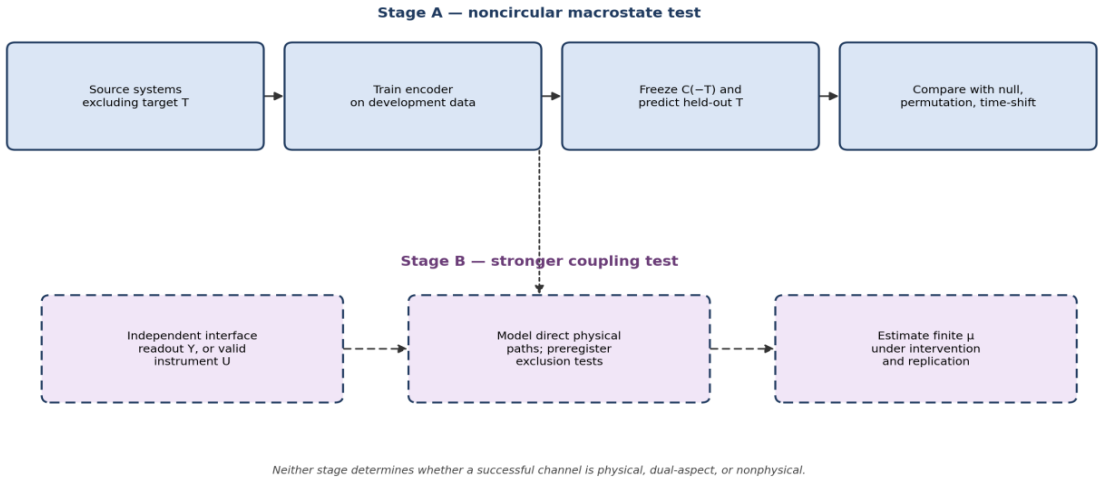

Figure 3. Stage A tests whether a cross-fitted macrostate improves prediction without target leakage. Stage B requires an independent interface readout or valid instrument before a coupling channel can be claimed. Neither stage determines whether a successful channel is physical, dual-aspect, or something else.

6.2 Does the method actually work? A reproducible test

A simulation shows the method can both succeed and correctly fail (full parameters in Appendix A.2, code in Appendix B, so every number here can be reproduced exactly). Three simulated physical subsystems followed Equation 2 with $\rho = 0.75$ and a true coupling of $\mu = (0.30, -0.22, 0.18)$; an independent interface readout $Y(t) = C(t) + \nu(t)$ was generated separately from the target outcomes. Trained on 1,750 of 2,499 time steps and tested on the remaining 749, adding Y cut mean held-out prediction error by 13.3% (Table 3) and recovered μ as approximately $(0.30, -0.21, 0.18)$ — close to the true values. Independently shuffling the Y sequences before refitting erased essentially all of that gain, confirming the improvement wasn't a statistical artifact. A second simulation with true $\mu = (0, 0, 0)$ — no real coupling — correctly showed no improvement from adding Y (a −0.04% change).

Target	H0 error	H1 error	Improvement	Estimated μ
B1	0.228	0.184	19.6%	0.297
B2	0.203	0.182	10.3%	−0.212
B3	0.210	0.190	9.4%	0.182
Mean	0.214	0.185	13.3%	—

Table 3. Held-out results for the synthetic coupling example.

This demonstrates identifiability under ideal conditions; it doesn't demonstrate that an independent interface readout exists in real biology. Its purpose is narrower: to replace an open-ended coupling placeholder with a finite model that can return a null result, and to give readers a script to run rather than numbers to take on faith. Every term in Equations 1–3 is implemented in Appendix B exactly as specified, with nothing added or silently dropped.

6.3 Raising the bar: a nonlinear competitor

Equation 1 is linear, and §6.2's simulation was generated from that same linear-Gaussian structure — a fair test of the inference procedure itself, but not a test against a physically richer competitor. Real physiology

is generally nonlinear, so a linear H0 can underfit, and part of any apparent gain from adding μC(t) could just be H1 soaking up nonlinear structure that was already there, with nothing to do with genuine macrostate coupling. Before treating a positive Tier 3 result as evidence for BMIC, it should be pitted against a preregistered nonlinear physical null instead — a nonlinear autoregressive or recurrent state-space model fit to B(t) alone, using the same held-out, cross-fitted, and permutation-controlled procedure from §6.1. A coupling gain that survives this harder comparison is much stronger evidence than one measured only against a linear baseline. A gain that doesn't survive just means the linear null was too weak a foil — not that BMIC is supported.

7. Relation to Competing Frameworks

7.1 Autopoiesis, constraint, and homeostasis

Autopoiesis explains living unity through organizational closure and self-production [3]; Deacon offers a constraint-based account of biological organization [7]; Damasio links feeling and mind to homeostatic regulation [8]. OIC sits comfortably in this tradition — it also treats unity as maintained organization. BMIC goes one step further, claiming that a distinct global representational state adds predictive or causal power beyond that organization. If closure and homeostasis turn out to explain everything on their own, BMIC simply isn't needed.

7.2 Active inference, causal emergence, and basal cognition

Active inference is BMIC's closest operational competitor — it already offers generative models, hierarchical state estimation, action, and viability-preserving dynamics [20,21]. Developmental bioelectricity treats morphogenesis as a multiscale, pre-neural problem-solving process, and identifies manipulable mechanisms that scale cellular goals up into anatomical outcomes [9–12]. Causal-emergence frameworks quantify when a macro-level description has real predictive or interventional advantages over its components [22,23]; the same effective-information logic could, in future work, offer an alternative way to estimate causal reach (c^e) beyond the intervention-based approach in Table 2. All three traditions reduce how much explanatory work is left for any additional macrostate to do — BMIC has to outperform them, not just resemble them.

8. Scope, and What the Companion Paper Adds

OIC makes no claim to logical or computational completeness — it's a bounded, functional measure of coordination, the same kind of bounded measure that's common and useful throughout biology and engineering. Nothing here turns on resolving deeper questions about the computational status of physical systems.

This paper has deliberately not asked whether the coordination it models bears on subjective experience in organisms with a nervous system. That's a real question, and a companion paper [25] takes it up directly: it imports the coupling law of Equations 1–3 unchanged and asks what would have to be added — specifically, a further, self-referential construct built on top of C(t), not folded into it — before the resulting state could be a candidate for anything experiential, and what evidential standard that claim should be held to. Neither paper's results are needed to evaluate the other, and nothing here should be read as having already answered that further question.

9. Predictions and How This Could Fail

The research program starts with physical measurement and ends, conditionally, with any claim about the candidate state's ontological status. A residual effect only counts as positive evidence if the model,

measurement channel, controls, and primary outcomes were all specified in advance — Table 4 lays out each test alongside the result that would weaken or falsify it.

Test	Required observation	Weakening or falsifying result
OIC construct validity	Component scores are reliable, converge with related measures, and prospectively predict recovery or state transitions.	Scores are unstable, redundant with a simpler metric, or fail to predict out of sample.
Cross-fitted macrostate	C(−T) improves held-out prediction of target T across rotated targets, independent cohorts, and a nonlinear physical null (§6.3).	The gain disappears with target exclusion, subject separation, time-shifting, surrogate controls, or a properly specified nonlinear competitor.
Finite coupling μ	μ reliably differs from zero, with a stable sign and magnitude, under intervention and replication.	Estimates shrink to zero, flip sign unpredictably, or depend on which flexible model is used.
Pre-neural interface	A specified early signal adds causal information beyond measured bioelectric, mechanical, and molecular pathways.	Expanded physical models and interventions fully account for developmental coordination without it.
Viability-recovery gradient	c^V declines gradually with perturbation severity, predicting partial as well as total loss of recovery (§3.3).	Recovery is effectively all-or-none, or fully explained by single-organ failure with no reference to cross-subsystem integration.
Ontological status	Operational results discriminate among physical, dual-aspect, and other candidate models of what the macrostate is.	Successful results remain equally well explained by physical latent-state or field models with no added ontological commitment.

Table 4. Preregistered tests and explicit weakening conditions.

The most informative first applications are narrow, high-dimensional systems where measurement and controlled perturbation can happen together. Developmental organoids and regenerative organisms can address the pre-neural claims; anesthesia transitions and intracranial or peripheral stimulation can address mature integration. These two programs should be analyzed separately — any claim of a single lifelong substrate would need to wait until both had been.

10. Proposed Empirical Program

10.1 Cross-subsystem integration study

A first study should combine measures from at least three of the five physical subsystems in Figure 1 — say, bioelectric, endocrine, and immune signals — in the same organisms. Source subsystems would estimate C(−T), with each target subsystem held out in turn, and the primary endpoint would be preregistered improvement in held-out prediction relative to subsystem-specific, whole-system latent-state, and nonlinear physical-null controls (§6.3). A robust cross-subsystem state would support Tier 2 or 3 functionality — it wouldn't settle the macrostate's ontological status.

10.2 Developmental and longitudinal refinement

Whether integrative capacity changes over the life course should be studied, not assumed. Age-appropriate measures of boundary participation, representation, causal reach, and viability recovery can show which OIC components improve, plateau, or reorganize with development — and whether they move together or

along separate trajectories. There's no reason to expect development to be monotonic across every component.

11. Discussion

OIC can be evaluated independently of BMIC, and might prove useful even if the stronger hypothesis fails outright. BMIC, in turn, can be tested as a candidate macrostate without presupposing anything beyond ordinary physical causation. Table 2's five components give OIC an explicit, falsifiable operationalization, rather than leaving "organismic wholeness" as an intuitive label that resists testing.

The single most important methodological commitment here is that Equations 1–3 are the only definition of C(t) used anywhere in this paper, and that the synthetic validation and the code in Appendix B implement exactly those equations — nothing added, nothing silently dropped. A framework that gets revised and extended over time is only as trustworthy as its willingness to keep its worked examples and its formal claims in exact correspondence; this paper treats that correspondence as a requirement, not a courtesy.

The central open empirical question is whether BMIC explains anything beyond active inference, network physiology, developmental bioelectricity, and causal emergence. That requires preregistered empirical superiority over both linear and nonlinear physical competitors (§6.3). If a stable global state or an independent coupling channel never turns up, that outcome supports the physical null model and requires withdrawing the stronger claim — a live and acceptable result, not a failure of the research program.

12. Conclusion

Organismic unity can be studied without assuming a hidden controller. OIC offers an explicit, bounded measure of boundary participation, representation, causal reach, temporal continuity, and recovery. BMIC adds a testable global-state hypothesis, with a single, consistently specified coupling law and a noncircular inference design — both fully reproducible from the parameters and code in Appendices A and B.

Current evidence supports distributed biological organization and distributed physiological dependence of organism-level coordination. It does not establish an ontologically distinct integrative substrate. BMIC remains a falsifiable research program whose future depends on whether its parameters survive target exclusion, active controls, nonlinear competitors, intervention, negative controls, and independent replication — and on whether any successful result can tell physical models apart from non-physical ones. A companion paper [25] takes up the further, separately argued question of whether any of this bears on self-reference or subjective experience.

References

1. Turing, A. M. (1952). The chemical basis of morphogenesis. Philosophical Transactions of the Royal Society B, 237, 37–72. https://doi.org/10.1098/rstb.1952.0012

2. Prigogine, I. (1978). Time, structure, and fluctuations. Science, 201, 777–785. https://doi.org/10.1126/science.201.4358.777

3. Varela, F. G., Maturana, H. R., & Uribe, R. (1974). Autopoiesis: The organization of living systems, its characterization and a model. BioSystems, 5, 187–196. https://doi.org/10.1016/0303-2647(74)90031-8

4. Cannon, W. B. (1929). Organization for physiological homeostasis. Physiological Reviews, 9, 399–431. https://doi.org/10.1152/physrev.1929.9.3.399

5. Bashan, A., Bartsch, R. P., Kantelhardt, J. W., Havlin, S., & Ivanov, P. C. (2012). Network physiology reveals relations between network topology and physiological function. Nature Communications, 3, 702. https://doi.org/10.1038/ncomms1705

6. Bartsch, R. P., Liu, K. K. L., Bashan, A., & Ivanov, P. C. (2015). Network physiology: How organ systems dynamically interact. PLOS ONE, 10, e0142143. https://doi.org/10.1371/journal.pone.0142143

7. Deacon, T. W. (2011). Incomplete Nature: How Mind Emerged from Matter. W. W. Norton.

8. Damasio, A. (2018). The Strange Order of Things: Life, Feeling, and the Making of Cultures. Pantheon.

9. Levin, M. (2022). Technological approach to mind everywhere: An experimentally grounded framework for understanding diverse bodies and minds. Frontiers in Systems Neuroscience, 16, 768201. https://doi.org/10.3389/fnsys.2022.768201

10. Pai, V. P., Aw, S., Shomrat, T., Lemire, J. M., & Levin, M. (2012). Transmembrane voltage potential controls embryonic eye patterning in Xenopus laevis. Development, 139, 313–323. https://doi.org/10.1242/dev.073759

11. Ciaunica, A., Levin, M., Rosas, F. E., & Friston, K. (2023). Nested selves: Self-organization and shared Markov blankets in prenatal development in humans. Topics in Cognitive Science. https://doi.org/10.1111/tops.12717

12. Manicka, S., Pai, V. P., & Levin, M. (2023). Information integration during bioelectric regulation of morphogenesis of the embryonic frog brain. iScience, 26, 108398. https://doi.org/10.1016/j.isci.2023.108398

13. Deglincerti, A., Croft, G. F., Pietila, L. N., et al. (2016). Self-organization of the in vitro attached human embryo. Nature, 533, 251–254. https://doi.org/10.1038/nature17948

14. Shahbazi, M. N., Jedrusik, A., Vuoristo, S., et al. (2016). Self-organization of the human embryo in the absence of maternal tissues. Nature Cell Biology, 18, 700–708. https://doi.org/10.1038/ncb3347

15. Yuan, S., Lin, H., Wu, C., et al. (2023). Human zygotic genome activation is initiated from the paternal genome. Cell Discovery, 9, 13. https://doi.org/10.1038/s41421-022-00494-z

16. Cao, J., O'Day, D. R., Pliner, H. A., et al. (2020). A human cell atlas of fetal gene expression. Science, 370, eaba7721. https://doi.org/10.1126/science.aba7721

17. Domcke, S., Hill, A. J., Daza, R. M., et al. (2020). A human cell atlas of fetal chromatin accessibility. Science, 370, eaba7612. https://doi.org/10.1126/science.aba7612

18. Greenbaum, S., Averbukh, I., Soon, E., et al. (2023). A spatially resolved timeline of the human maternal–fetal interface. Nature, 619, 595–605. https://doi.org/10.1038/s41586-023-06298-9

19. Caldarelli, P., Chamolly, A., Villedieu, A., et al. (2024). Self-organized tissue mechanics underlie embryonic regulation. Nature, 633, 887–894. https://doi.org/10.1038/s41586-024-07934-8

20. Friston, K. (2010). The free-energy principle: A unified brain theory? Nature Reviews Neuroscience, 11, 127–138. https://doi.org/10.1038/nrn2787

21. Parr, T., Pezzulo, G., & Friston, K. J. (2022). Active Inference: The Free Energy Principle in Mind, Brain, and Behavior. MIT Press.

22. Hoel, E. P., Albantakis, L., & Tononi, G. (2013). Quantifying causal emergence shows that macro can beat micro. Proceedings of the National Academy of Sciences, 110, 19790–19795. https://doi.org/10.1073/pnas.1314922110

23. Rosas, F. E., Mediano, P. A. M., Jensen, H. J., Seth, A. K., Barrett, A. B., Carhart-Harris, R. L., & Bor, D. (2020). Reconciling emergences: An information-theoretic approach to identify causal emergence in multivariate data. PLOS Computational Biology, 16, e1008289. https://doi.org/10.1371/journal.pcbi.1008289

24. Dennett, D. C. (1991). Consciousness Explained. Little, Brown and Co.

25. Moganti, M. The Self That Reports: Self-Referential Integration and the Meta-Problem of Consciousness. Companion paper, in preparation.

Appendix A. Technical Details for the Worked Examples

Both worked examples in the main text are fully reproducible from the parameters below; the complete analysis script is in Appendix B.

A.1 OIC worked example (§5.2)

Component scores were assigned by hand to illustrate index behavior, not measured from data: integrated state (c^B, c^R, c^c, c^T, c^V) = (0.88, 0.82, 0.76, 0.84, 0.90); perturbed state = (0.74, 0.60, 0.35, 0.50, 0.42). With equal weights ($w_j = 0.2$), the geometric mean (Equation 4) gives OIC = 0.839 for the integrated state and 0.504 for the perturbed state.

A.2 Coupling worked example (§6.2)

The simulation used T = 2,500 time points and three physical subsystems following Equation 2, with random seeds fixed at 42 (nonzero-coupling run) and 43 (zero-coupling control):

$A = [[0.55, 0.10, 0.00], [0.05, 0.60, 0.08], [0.00, 0.12, 0.50]]$

$\mu = (0.30, -0.22, 0.18)$ for the coupling run; $\mu = (0, 0, 0)$ for the null control

$\rho = 0.75$; $C(t+1) = \rho \cdot C(t) + 0.35 \cdot U(t-1) + \eta(t)$, with U(t) an independent standard-normal source signal

$Y(t) = C(t) + \nu(t)$

Process-noise standard deviations were 0.18 for each B_i, 0.20 for C, and 0.10 for Y; all stochastic terms were independent, zero-mean, and Gaussian. Ordinary least squares was fit to the first 1,750 of 2,499 transitions and evaluated on the remaining 749. H0 used B(t−1) alone to predict B(t); H1 added Y(t−1). The permutation control independently shuffled the training- and test-set Y sequences before refitting H1. Appendix B reproduces every value in Table 3 exactly, given these seeds.

Appendix B. Analysis Code

The script below (Python 3, NumPy only) generates both simulations from Appendix A.2 and reproduces every number in Table 3, including the permutation and null-coupling controls. It implements Equations 1–3 exactly, with nothing added or omitted, so a reader can run it directly rather than take the reported values on faith.

```
import numpy as np

def run(true_mu, seed):
    rng = np.random.default_rng(seed)
    T = 2500
    A = np.array([[0.55, 0.10, 0.00],
                  [0.05, 0.60, 0.08],
                  [0.00, 0.12, 0.50]])
    mu, rho = np.array(true_mu), 0.75
    sd_B, sd_C, sd_Y = 0.18, 0.20, 0.10
    B, C = np.zeros((T, 3)), np.zeros(T)
    U = rng.normal(0, 1.0, T)
    B[0] = rng.normal(0, 1, 3)
    C[0] = rng.normal(0, 1)
    for t in range(1, T):
        C[t] = rho * C[t-1] + 0.35 * U[t-1] + rng.normal(0, sd_C)
        B[t] = A @ B[t-1] + mu * C[t-1] + rng.normal(0, sd_B, 3)
    Y = C + rng.normal(0, sd_Y, T)
    n_tr = 1750
    X0_tr, X0_te = B[:n_tr-1], B[n_tr-1:-1]
    Y_tr, Y_te = Y[:n_tr-1], Y[n_tr-1:-1]
    tgt_tr, tgt_te = B[1:n_tr], B[n_tr:]
    A_hat, *_ = np.linalg.lstsq(X0_tr, tgt_tr, rcond=None) # H0
    rmse0 = np.sqrt(np.mean((tgt_te - X0_te @ A_hat) ** 2, axis=0))
    X1_tr = np.column_stack([X0_tr, Y_tr]) # H1
    X1_te = np.column_stack([X0_te, Y_te])
    coef1, *_ = np.linalg.lstsq(X1_tr, tgt_tr, rcond=None)
    mu_hat = coef1[3]
```

```
    rmse1 = np.sqrt(np.mean((tgt_te - X1_te @ coef1) ** 2, axis=0))
    pr = np.random.default_rng(seed + 1000) # permutation control
    Y_tr_p, Y_te_p = pr.permutation(Y_tr), pr.permutation(Y_te)
    X1p_tr = np.column_stack([X0_tr, Y_tr_p])
    X1p_te = np.column_stack([X0_te, Y_te_p])
    coefp, *_ = np.linalg.lstsq(X1p_tr, tgt_tr, rcond=None)
    rmsep = np.sqrt(np.mean((tgt_te - X1p_te @ coefp) ** 2, axis=0))
    return rmse0, rmse1, rmsep, mu_hat

rmse0, rmse1, rmsep, mu_hat = run([0.30, -0.22, 0.18], seed=42)
print("Coupling run H0:", rmse0.round(3), " H1:", rmse1.round(3),
      " permuted-Y:", rmsep.round(3), " mu_hat:", mu_hat.round(3))
rmse0b, rmse1b, *_ = run([0.0, 0.0, 0.0], seed=43)
print("Null control H0:", rmse0b.round(3), " H1:", rmse1b.round(3))
```