systems

MDPI

Article

The Coercivity Law of Enaction Within Fisher-Generative Informational Realism: A Cybernetic Threshold for Autopoietic Closure

Maurice Yolles [1,*] and Chin-Ken Lin [2]

1 Liverpool Business School, Liverpool John Moores University, Liverpool L3 5UG, UK
2 Independent Researcher, New Taipei City 251, Taiwan; kingkonglin2050@gmail.com
* Correspondence: prof.m.yolles@gmail.com

Highlights

Please indicate how your work links to systems science via your contributions to systems practice, theory, and/or methodology.

- To become autopoietic, a complex adaptive system requires both a structure capable of third-order closure and sufficient resources to sustain it under stress. That requirement is proto-energetic rather than physical: proto-energy is the informational overhead, measured on the system's Fisher-information geometry, needed to hold its shape, rules, and memory steady against perturbation—an analogue of the continuous throughput by which far-from-equilibrium order is maintained. Deriving this constraint as a coercivity law gives Ashby's Law of Requisite Variety a "cost side": variety must be affordable on the state manifold, not merely sufficient in quantity.
- Offers a practical way to estimate this proto-energetic condition in real organizations, using proxies for how much regulatory information a system holds and how quickly it locks in under pressure, checked against a resonance-signature test rather than a physical calibration constant unavailable outside physics.

What are the main findings and/or the implications of the main findings?

- Derives a coercivity threshold marking the point where a system's informational structure becomes self-sustaining rather than dependent on outside support. Connecting this threshold to any real-world system, physical, biological, or organizational, depends on conditions specific to that system's context, which the paper identifies but does not prove, so the threshold is not presented as a ready-made law for any one domain.
- Classifies large language models as proto-autopoietic: under external scaffolding they display transient structures resembling third-order closure, but withdraw the scaffolding, and these dissipate within a single inference call, since no persistent internal field accumulates the rigidity that self-production requires. Crossing the coercivity threshold would require persistent memory or a persistent self-model, not more parameters.

Abstract

When does a complex adaptive system (CAS) (like a thermostat, a market, a flock, or a large language model) become a complex adaptive autopoietic system (CAAS), one that actively produces and sustains itself rather than merely adapting to its environment? We argue that this transition requires two simultaneous conditions. The first, established in a companion work, is architectural sufficiency. The system's decision structure must achieve recursive

Check for updates

Academic Editor: Gianfranco Minati

Received: 20 July 2026
Revised: 29 August 2026
Accepted: 3 September 2026
Published: 11 September 2026

closure at the third cybernetic order, the so-called fractal-seed point. The second, derived here, is corporeal viability. The system must pay a structural cost (the enactment tension E) large enough to hold its organization in place against perturbation. Working within Fisher-Generative Informational Realism (FGIR), an informational-realist framework in which information, not matter or energy, is the primary generative substrate of reality, we derive a scalar invariant $E = I_{\mathrm{p}}\mathrm{c}^2$, where I_{p} is the integrated informational potential of the operative field and c is the global coherence conductance at which informational structure locks into place. This invariant marks the threshold at which a proto-autopoietic system crosses into full autopoiesis. Because FGIR treats physical mass-energy as the result of a freezing projection acting on an incorporeal informational manifold rather than as the foundation of reality, the same invariant that governs autopoietic closure in social and organizational systems also projects, under appropriate bridge conditions, for instance, onto the classical physical law $\hat{E} = \hat{m}\,\hat{c}^2$, or in a quantum context, the Schrödinger equation. The physical projection is stated as a conditional equivalence (T-BRIDGE), but it is not the focus of this paper. Rather, the central contribution is the cybernetic threshold itself and its operationalization. We show how a systems theorist can assess a system's distance from critical admissibility, even when the absolute magnitudes of I_{p} and c are unavailable in domains that lack R(6) closure; we connect the coercivity threshold explicitly to Ashby's Law of Requisite Variety, supplying the energetic constraint that Ashby's purely combinatorial criterion leaves implicit. The contemporary case of large language models illustrates the framework. LLMs display transient, externally-scaffolded R(3)-like decision structure but fail the coercivity threshold, and so remain proto-autopoietic. The paper thus offers systems science a derived criterion, additional to architectural closure, for distinguishing full from proto-autopoiesis, with a specified but as yet unexecuted test program across physical, biological, cognitive, and social domains.

Keywords: cost of persistence; informational realism; complex adaptive systems; autopoiesis; cybernetics; requisite variety; large language models; systems theory

1. Introduction

Fisher-Generative Informational Realism (FGIR) emerges from metacybernetic theory [1], the framework that formalizes how reflexive systems generate their own governing principles across successive cybernetic orders. Living systems do not passively sense a ready-made world. They actively bring their world into being through their own organization, a process the biologists Varela, Thompson and Rosch [2] called enaction. Eric Schwarz's third-order cybernetics supplied a systemic architecture for this insight: autogenesis, the recursive closure by which viable systems self-organize and embed their own governing principles at a boundary, escaping infinite regress [3–6]. Both traditions, however, left a gap. Neither supplied a quantitative measure of the structural effort required to convert a merely dispositional, potential state into a stable, enacted one. This paper provides that measure and, in doing so, offers systems science a quantitative threshold at which a complex adaptive system becomes a self-producing one.

The question we address is squarely cybernetic. Holland [7,8] defined a complex adaptive system (CAS) as adaptive, pattern-forming, and rule-evolving, but not necessarily self-producing. A complex adaptive autopoietic system (CAAS) additionally possesses full autopoietic closure in Maturana and Varela's sense [9,10]: both operative closure (moment-to-moment self-maintenance, contingent on external boundary conditions) and dispositional closure (a self-legislating feedback loop that regenerates the system's own

https://doi.org/10.3390/systems14091132

governing parameters). What distinguishes the two? Existing cybernetic theory offers architectural criteria (recursion depth, requisite variety, viable-system structure), but no quantitative criterion for the energetic or structural cost a system must pay to cross from CAS to CAAS. We argue that two conditions are jointly necessary, and we derive the second of them here.

Our central thesis is that the CAS-to-CAAS transition requires two conditions to be met simultaneously. The first, architectural sufficiency, is the R(3) fractal-seed closure established in the companion metacybernetics paper [1]. That paper derived the recursion point at which a system's decision boundary first acts on outputs that structurally mirror its own inputs, producing the canonical 2–3–2–3 parity alternation of trait cardinality across the cybernetic-order hierarchy R(1), R(2), R(3), . . . It established where recursive closure occurs and why it takes the trait-structure it does; it did not ask what closure costs. R(3) closure is architecturally primary (a precondition without which autopoiesis is structurally impossible), but it is not by itself sufficient. Two systems can share the same decision-boundary topology, the same 2–3–2–3 trait parity, and the same formal closure of the feedback loop, and yet behave very differently under stress: one recovers its organizational identity after a severe perturbation, and the other does not.

The second condition, corporeal viability, is the subject of this paper. We derive a scalar invariant, the enactment tension E, that quantifies the minimum rigidity a system must generate to transition from proto-autopoietic viability to full, self-producing autopoiesis. Proto-autopoiesis denotes a structurally incomplete mode of operation in which a system can sustain its operative patterns but cannot stabilize or regenerate its own governing parameters without external scaffolding. Viruses, as analysed in 2022 by Yolles & Frieden [11], exemplify this condition: they qualify as living systems in their organizational and informational structure, yet they do so in a proto-autopoietic mode, enacting operative autopoiesis while relying on host-provided boundary conditions to maintain viability and reproduce their governing parameters.

The invariant takes the form $E = I_\mathrm{p}c^2$, where I_p is the integrated informational potential of the system's operative field and c is the global coherence conductance at which informational structure locks into place. A system with R(3) architectural closure but sub-threshold $I_\mathrm{p}c^2$ has the right shape for autopoiesis that has not yet paid the rigidity cost that would make it actually persist against perturbation. The Coercivity Law, therefore, completes, on the energetic axis, what the fractal-seed paper established on the architectural axis. Neither condition alone suffices.

This distinction matters because the systems-science literature, from Holland onwards, has long recognized that CASs can exhibit rich, rule-based, self-organizing behavior (flocking, market dynamics, immune response) without any of that behavior being self-producing in the strict sense of Maturana and Varela. What has been missing is a quantitative axis on which that distinction can be drawn. The coercivity law supplies one. A system that loses dispositional closure (whose feedback loop stops regenerating its own governing parameters) does not instantaneously stop functioning. It degrades into proto-autopoietic operation, sustaining its immediate operative patterns for as long as the environment continues to supply the boundary conditions it would otherwise supply internally. Organizationally, this is the condition of an institution that continues to execute its standard procedures while its capacity to revise, defend, or regenerate the values underlying those procedures has been captured or hollowed out by an external actor. The gap between a system's actual $I_\mathrm{p}c^2$ and the coercivity threshold is, in principle, a measurable index of how far a system sits from autonomous self-production, rather than a binary label attached after the fact.

The two-condition framework gives a precise formal vocabulary for a question that has become pressing in contemporary systems theory: are large language models (LLMs) autonomous agents, or merely pattern-completers? We argue that LLMs are best classified as proto-autopoietic. They display, over short horizons and under external scaffolding, decision structures that resemble R(3) closure: they can be prompted to monitor and revise their own outputs (R(2)-type self-stabilization), and chain-of-thought or agentic scaffolds can elicit something that looks like regulation of the conditions of their own stability (R(3)-type behavior). Architecturally, then, an LLM-in-its-scaffolding exhibits a transient, externally-supported fractal-seed-like structure.

What LLMs lack is the second condition. They do not pay the structural cost required to persist independently of their scaffolding. Their operational coherence depends on continued training, platform engagement, and prompting, all of which supply, externally, the boundary conditions the system's own K-field would otherwise have to supply internally. Withdraw the scaffolding, and the R(3)-like decision structure dissipates within a single inference call: there is no accumulated rigidity on the operative manifold that would let the system regenerate its own governing parameters. In the vocabulary of this paper, LLMs sit below the coercivity threshold. They are proto-autopoietic in exactly the same sense as a Bénard convection cell, operationally closed while the boundary conditions hold, but incapable of dispositional closure. This is not a deficit to be remedied by scale; it is a structural fact about where the cost of persistence is being paid. The framework predicts that no amount of parameter scaling alone will cross the coercivity threshold, because the threshold is a property of the system's operative manifold, not of its parameter count. This example is not incidental. The cybernetic literature has, since Ashby, distinguished systems that maintain themselves from systems that merely respond. The contribution of the coercivity law is to make that distinction quantitative and to locate contemporary AI systems precisely on the resulting axis. We return to the operational details (how one would actually estimate I_p and c for an LLM or for an organization) in Section 8.

The Fisher-Generative Informational Realism (FGIR) framework within which we work is an informational-realist ontology. It treats information, not matter or energy, as the primary generative substrate of reality. What we ordinarily call physical objects are stabilized patterns of information that have undergone enaction. Because the framework is informational-realist, physical mass-energy is not the foundation of reality but the result of a freezing projection π acting on an incorporeal informational manifold. The same informational invariant $E = I_\mathrm{p}c^2$ that we derive here as the threshold of autopoietic closure in social and organizational systems, also projects, under appropriate bridge conditions, onto the physical law of special relativity $\hat{E} = \hat{m}\,\hat{c}^2$. This physical projection forms a conditional equivalence (the T-BRIDGE theorem of Section 7) which establishes the ontological grounding of the framework, though the central contribution of this paper is the cybernetic threshold itself.

This ordering matters for a systems-science readership. The framework's claim is not that social systems are physics-like, nor that organizational behavior reduces to mass-energy. The claim runs in the other direction. It is that physical mass-energy is one domain, and hence one projection, of a more general informational invariant. The same invariant also governs the persistence of biological, cognitive, and social systems. The coercivity law is a theorem about information manifolds. Its physical reading is one consequence among several. Readers from physics will find the bridge to $\hat{E} = \hat{m}\,\hat{c}^2$ in Sections 6 and 7. Readers from cybernetics and systems science can proceed directly from Sections 1–4 and 8, where the threshold is derived and applied without requiring the physical bridge.

FGIR's ontology comprises four fields. Three are incorporeal, and in a third-order cybernetic system indicated by the recursive symbol (3), they are the dispositional field J (latent intentions and rules), the operative field I (moment-to-moment processing), and the sustentative field K (slow-moving background context). These constitute an informational manifold that contains no energy, mass, or spacetime. The fourth, the corporeal field C, is where physical and observable structure resides, and is where the freezing projection π delivers its output. This four-field architecture was consolidated by the Fisher Information Field Theory (FIFT) program [12] and refined by subsequent work on Fisher-metric tracking and ontological projection [13], quantum-informed cybernetics [14], and empirical applications to civilizational and geopolitical dynamics [15,16]. Section 2 develops the ontology in more detail; for now, the key point is that the enactment tension E is a purely informational quantity (defined on the incorporeal manifold) that measures how much a system's space of possibilities must be constrained before a configuration can lock in and persist. It is one level upstream of the physical world.

Readers whose interest is systemic and cybernetic, rather than differential-geometric, can follow the argument's logic without the equations. Three ideas do all the conceptual work. First, every persisting system has to pay for its own persistence. A pattern that merely could exist is cheaper, ontologically speaking, than a pattern that actually does exist and keeps existing under perturbation. Second, that cost is geometric before it is physical: it is a measure of how much a system's space of possibilities has to be bent or constrained for one configuration to lock in, and this bending can be measured on the abstract space of informational possibilities before any energy, mass, or spacetime is in the picture at all. Third, the same measure, read in a physical context, is energy: what looks like informational rigidity from the inside of the manifold looks like mass-energy once the system is embedded in ordinary spacetime. The mathematics in Sections 3–7 is the precise version of these three sentences; a reader who accepts them provisionally can proceed directly to Section 8, where the payoff (a table of what this cost looks like across a dozen different kinds of systems, from photons to organizations, together with operational guidance on how to estimate it) is stated with a minimum of formal apparatus. The heavy differential-geometric detail (Bakry–Émery metrics, Chentsov's theorem, Picard–Lefschetz monodromy) is collected in Appendix B for readers who require it; the main text requires only the logical skeleton of the proof.

The paper makes three contributions. First, it offers a formal derivation, from a Saturation-Rigidity Principle and the geometry of weighted information manifolds, of the enactment-tension invariant $\mathrm{E} = \mathrm{Ipc}^2$, rather than its assumption. Second, it provides a quantitative criterion (additional to the architectural R(3)-closure criterion of [1]) for distinguishing full from proto-autopoiesis in complex adaptive systems, together with a dimensionless saturation deficit by which social and organizational systems may be assessed for distance from critical admissibility, absolute magnitudes being unavailable there (Section 8). Third, it records a conditional bridge theorem (T-BRIDGE) linking the incorporeal invariant to its corporeal physical reading. Its three bridge conditions are undischarged, so it is recorded as an open program rather than offered as a result of this paper. Throughout, claims are hedged using standard academic conventions: results that are proven are stated as theorems; results that depend on adopted conditions are stated as conditional; results that await empirical test are identified as conjectural. Section 2 introduces the FGIR ontology. Section 3 states the Saturation-Rigidity Principle. Section 4 derives the Fundamental Law $\mathrm{E} = \mathcal{F}c^2$ on a general information manifold, with technical detail deferred to Appendix B. Section 5 specializes it to $\mathrm{E} = \mathrm{Ipc}^2$ via the Fisher–Rao coarse-graining structure. Section 6 formalizes the freezing projection π and the three bridge conditions. Section 7 states the T-BRIDGE conditional equivalence theorem. Section 8 (the central section for systems-theory readers) catalogues the corporeal projections and operationalizes the cost: how to estimate I_p and c in an organization, how the framework connects to Ashby's Law of Requisite Variety, and how the LLM example is formalized. Section 9 gives a consolidated derivation summary, and Section 10 closes with the research program and the epistemic-tier table.

To assist in the reading of this article, it is useful to list the key terms used, as shown in Table 1, while a full notation glossary, cross-referenced to its point of introduction, is given in Appendix A.

Table 1. Notational Glossary.

Term	Meaning
CAS/CAAS	Complex Adaptive System/Complex Adaptive Autopoietic System. A CAS adapts and forms patterns; a CAAS additionally self-produces. The paper's central question is when the former becomes the latter.
FGIR	Fisher-Generative Informational Realism. The framework within which the derivation is conducted treats information, not matter or energy, as the primary generative substrate of reality.
Four fields (J, I, K, C)	Dispositional (J: latent possibilities), operative (I: moment-to-moment processing), sustentative (K: long-horizon context), and corporeal (C: observable outputs). The first three are incorporeal; C is where physical observables reside.
R(n)	The n-th order of cybernetic recursion. R(3) is the fractal-seed point at which recursive closure first becomes structurally possible; R(6) is where the freezing projection is defined.
SRP	Saturation-Rigidity Principle. Three conditions (capacity saturation, defect rigidity, rapid locking) whose joint imposition produces constitutional persistence.
$E = I_p c^2$	The coercivity law. E is the enactment tension (structural cost of persistence); Ip is the integrated informational potential; c is the coherence conductance. This is the threshold a system must cross to become autopoietic.
π (freezing projection)	The operator that reassigns an informational quantity from the incorporeal manifold to the corporeal field, where it acquires physical meaning.
T-BRIDGE	The conditional equivalence theorem: under bridge conditions BC1–BC3, $E = I_p c^2$ and $\hat{E} = \hat{m}\hat{c}^2$ are the same equation in two ontological registers.
CPC	Constitutional Persistence Cycle. A three-step construction (rigidity lock, entity collapse, monodromy reset) that generates the telegrapher's equation as the common ancestor of the Schrödinger, diffusion, and Dirac equations.
$T^E{}_0/T^E{}_1/T^E{}_2$	Epistemic tiers: $T^E{}_0$ = proven; $T^E{}_1$ = conditional/sketch-proven; $T^E{}_2$ = conjectural/awaiting empirical test.

2. The Ontological Framework of FGIR

Reality, in FGIR, is structured by three incorporeal informational fields whose geometric interactions generate the conditions under which corporeal structure can exist. The dispositional field J carries a system's latent possibilities, the rules, intentions, and value commitments from which concrete action is elaborated. The operative field I carries the moment-to-moment processing through which those possibilities are elaborated into behavior. The sustentative field K carries slow-moving boundary conditions and long-horizon context, which, in an organizational context, one would call culture and institutional memory. Together, J, I, and K constitute an informational manifold that contains no energy, mass, or spacetime; it is the substrate on which the enactment tension is defined. The corporeal field C is where physical observables reside (i.e., the enacted outputs, decisions, and artefacts an external observer can actually see), and this is the target of the freezing projection π. This four-field ontology descends from early metacybernetic work [17], was consolidated as the Fisher Information Field Theory in [12], and has been refined by subsequent work on Fisher-metric tracking [13] and quantum-informed cybernetics [14].

The reader coming from systems or management cybernetics may find a rough organizational analogy helpful, offered here only as an aid to intuition and not as part of the formal derivation. The sustentative K-field plays a role something like an organization's culture and institutional memory: slow-moving, providing continuity and background context rather than issuing moment-to-moment instructions. The dispositional J-field plays a role, something like strategy and value commitments. The operative I-field plays a role, something like execution: the concrete processing through which strategy becomes behavior. The corporeal C-field is observable, constituting the enacted outputs that an external observer can see. On this reading, the coercivity law claims that turning latent strategic intent into durable, self-sustaining organizational behavior is not free: it costs a rigidity that scales with how much operative information has to be locked into place, and organizations that have not paid this cost remain proto-autopoietic.

A closely related interpretation comes from evolutionary economics, where long-run persistence is understood in terms of routines, capabilities, and selection environments. In this register, the sustentative K-field corresponds to slow-moving institutional and technological regimes, the inherited structures that shape the selection landscape [18,19]. The dispositional J-field corresponds to search heuristics, innovation strategies, and expectation-formation [20], while the operative I-field corresponds to the routines and capabilities through which firms enact those strategies. The corporeal C-field corresponds to market-level observables (prices, output, innovation events, entry and exit) through which evolutionary selection acts. On this reading, the coercivity law formalizes a core evolutionary-economic insight: stabilizing a routine or capability requires a rigidity cost. Systems that have not paid this cost behave like proto-autopoietic economies, exhibiting transient adaptive behavior but unable to regenerate their own routines or maintain their identity when selection pressures change [21,22]. As shown in [22], metacybernetic analysis of market ideologies already treats capitalist systems as reflexive, adaptive, and selection-driven; evolutionary economics and FGIR therefore converge on the view that persistence is the stabilization of informational structure under selection.

Coupling between the four fields is realized by three orthogonal process intelligences, which it is useful to name here because they recur throughout the paper. Autopoiesis is the J$\leftrightarrow$I channel, the self-regenerating dispositional–operative couple, governed by the scalar ratio $\lambda = J/I$. Autogenesis is the K$\leftrightarrow$(J–I) channel, through which the sustentative field supplies long-horizon boundary conditions. Autopraxis is the (J–I–K)$\leftrightarrow$C channel, the locus of enacted practice where distributed informational potentials freeze into corporeal observables. This three-channel structure maps directly onto the R(1)–R(3) recursive hierarchy of [1], where autopoiesis is the dyadic R(1)-type relation, autogenesis the triadic R(2)-type relation, and autopraxis the R(3) fractal-seed closure of the incorporeal triad. Two structural distinctions matter for what follows. Each field has its own lateral implicate–explicate interface (the B-manifolds, internal to a single field), distinct from the transverse

H-manifold, which couples the incorporeal triad to the corporeal field across realms; the H-manifold's quantum-mechanical limit is ordinary quantum entanglement.

The freezing projection π is the operator that reassigns an informational quantity from the incorporeal manifold to the corporeal implicate field C_{impl}, where it acquires a concrete physical meaning. It applies whenever a system's informational structure is given a corporeal interpretation; what differs between contexts is only how the frozen quantity is read, as mass in a relativistic context, as intensity or charge density in an electromagnetic one, as concentration in a diffusive one, or as $|\psi|^2$ in a quantum-coherent one. Before freezing, the enactment tension Einc is a purely informational scalar defined on the incorporeal manifold; after freezing, the same number is read as a corporeal quantity Ecorp, whose concrete meaning depends on the local corporeal context. Freezing preserves the magnitude of the invariant while changing its ontological status: the number stays the same, but its meaning shifts from the rigidity required for information to stabilize itself to the energy that appears when that stabilized structure enters corporeality. Three everyday analogies make the idea concrete: surface tension becoming a rigid ice boundary on freezing; internal stress in a material becoming stored elastic energy on fixing; and potential-field curvature becoming a measurable force once a particle is placed in it. In each case, a pre-physical tension becomes a physical structure only once a rigidity threshold is crossed.

The four-field ontology is adopted here as a working commitment of the FGIR framework, grounded in the framework's axiomatic basis (Axioms A1–A10 of the FIFT Schema) rather than derived as a theorem. A1 asserts that information is ontologically primary; A2 asserts that the incorporeal fields J, I, and K are ontologically real and constitute the generative substrate of C; A7 asserts that third-order cybernetic closure is required, with K meta-regulating J and I without infinite regress; and A10 asserts that systemic existence is the maintenance of informational coherence above a minimal Fisher-information threshold. On these axioms, the four-field architecture is not an arbitrary modelling choice but a functional commitment: each field is justified by the operational role it plays in the SRP and in the freezing projection. A structural correspondence holds between the three SRP conditions of Section 3 and the three incorporeal fields (J with capacity saturation, I with defect rigidity, K with rapid locking), and the corporeal field C is the minimal image of the incorporeal triad under the freezing projection π. Whether this four-field structure is minimal is addressed in the Minimality Conjecture (Section 10), where a category-theoretic sketch shows that the interactions among J, I, and K generically require a fourth field C. The Minimality Conjecture concerns the size of the ontology required to support both third-order cybernetic closure and the freezing projection. Third-order cybernetics is defined entirely through the incorporeal triad of J, I, and K. These three fields are sufficient for recursive closure and for the fractal-seed architecture [1]. What enactment adds is the requirement that the informational structure have a well-defined corporeal image under the freezing projection π. The Minimality Conjecture states that, given the interaction structure of the triad J–I–K, a distinct corporeal field C is generically forced. It is the smallest additional field that allows the projection π to be defined coherently and allows third-order closure to be read in corporeal terms. In this sense, the conjecture is the formal analogue of the parsimony principle within FGIR, as explained in Yolles [23], which asserts that the four-field ontology is not an arbitrary elaboration but the smallest architecture capable of supporting both the fractal seed and its corporeal realization. If discharged, the Minimality Conjecture would convert the ontology from an axiomatic commitment into a derived consequence; until then, the four-field structure is treated as a functional commitment justified by its operational role in the framework.

https://doi.org/10.3390/systems14091132

3. The Saturation-Rigidity Principle

Before formally stating the Saturation-Rigidity Principle, it is useful to see how persistence can fail within the geometric framework developed in Sections 4 and 5. The admissible protoenergy functional (defined as the geometric cost of informational configuration on the incorporeal manifold) has three independent geometric features: a unique saturated configuration, a curvature-based penalty for deviation, and a gradient-driven return dynamic. Each feature corresponds to one of the SRP conditions, and the absence of any one of them produces a characteristic failure mode.

Firstly, a system may reach the richest configuration its constraints permit, the maximal Fisher information state admissible on its manifold, and yet remain fragile, since any small perturbation may dislodge it permanently, and nothing in its structure penalizes deviation. Secondly, a system may penalize deviation heavily, resisting perturbation strongly while sitting at that configuration, but be slow to return once perturbed, drifting back only asymptotically rather than locking in. Here, the curvature of the capacity landscape is present but insufficiently coupled to a rapid locking mechanism. Thirdly, a system may recover quickly from perturbation without ever having reached a genuinely saturated configuration in the first place, oscillating efficiently around a moving target rather than a fixed one.

Within FGIR, these three cases are exhaustive. The structure of the protoenergy functional ensures that any failure of persistence corresponds to the failure of at least one of the three geometric conditions. None of the systems described above persists in the strong sense that FGIR requires, in which a living cell, a stable institution, or a bound particle holds its organizational identity through time. Persistence, on this reading, is the conjunction of reaching the richest admissible configuration (saturation), resisting departure from it in proportion to the attempted deviation (rigidity, governed by local curvature), and returning to it faster than merely asymptotically once perturbed (rapid locking).

The Saturation–Rigidity Principle states that constitutional persistence arises only when three jointly necessary conditions hold. Each condition captures a distinct structural requirement, and none is sufficient alone. Their joint imposition produces a stable enacted configuration: a subsystem that attains maximal admissible structure, resists perturbation, restores itself with superlinear speed, and, when corporeally projected, acquires temporal and causal organization characteristic of Lorentzian domains. The three conditions are defined as follows:

- S1. Capacity Saturation. A subsystem attains the maximum information, energy, or complexity permitted by its constraints.
- S2. Defect Rigidity. Departures from saturation incur a cost that grows with local second-order structure:

  $$S(x) = \| \nabla_G W(x) \|_G^2,$$

  for the capacity functional W and Riemannian metric G.
- S3. Rapid Locking. Return to saturation is superlinearly fast, governed by a Łojasiewicz gradient inequality [24] near an isolated critical point x^*:

  $$| W(x) - W(x^*) |^{1-\alpha} \leq K_L \| \nabla_g W(x) \|_g, \alpha \in (0, 1/2]. \quad (1)$$

  When $\alpha > 0$, convergence accelerates as the defect becomes small, marking a phase transition in the return dynamics.

The three conditions above are the whole of the Saturation-Rigidity Principle. Nothing further is required to fix the causal character of the corporeal projection, which follows from the coercivity constant itself.

Proposition 1 (Cone structure from bounded coercivity)**.** *Let c be finite and non-zero, and let A^* be a saturation set on which the norm of ∇I in the admissible metric g equals c. Then $u := \frac{\nabla I}{c}$ is a unit vector field on A^*, determined by I and g alone, and the symmetric form $g_L := g - 2\,u^\flat \otimes u^\flat$ is Lorentzian, with u spanning its unique timelike direction. Its null cone defines the causal structure of g_L; the identification of that cone with the boundary of admissible information flow is the work of BC3. Since a non-degenerate quadratic form is determined by its null cone up to a positive factor, g_L is the unique such form up to positive conformal scaling.*

Proof. Finiteness of c makes the quotient ∇I/c defined, and gradient saturation makes it of unit norm, so u exists and vanishes nowhere on A^*. Writing a displacement as $v = au + w$ with w orthogonal to u gives $g(v,v) = a^2 + \parallel w \parallel^2$ and $g_L(v,v) = \parallel w \parallel^2 - a^2$, a form with one negative and $n-1$ positive eigenvalues whose null locus is the cone $a^2 = \parallel w \parallel^2$ with axis u. The uniqueness clause is the standard fact that a non-degenerate quadratic form is fixed by its zero set up to a positive scalar. □

Two things should be marked. The first is that the construction consumes exactly two hypotheses already in force: finiteness of c, required by the coercive upper bound of Section 4, and its non-vanishing, required by gradient saturation. We record the hypothesis explicitly, since it has been used implicitly throughout: $c \in (0, \infty)$. The second is what the construction does not supply. It fixes the signature and the timelike direction, not the dimension: the projection acquires one timelike and $n - 1$ spacelike directions for whatever n the underlying manifold carries, and that $n = 4$ is the residual content of the bridge condition BC1. Nor does it fix the scale relating the cone to a measured propagation speed, which is the work of BC3.

The condition previously stated at this point required the information gradient to be non-spacelike with respect to a fixed-point metric. Since the admissible geometries are Riemannian, and under a Riemannian metric every non-zero vector is spacelike, that condition was satisfiable only by the zero vector, which gradient saturation excludes. It was an error of formulation rather than of substance, and the result it was introduced to secure follows from the coercivity bound without it.

Now, S1 ensures that the subsystem occupies the extremal admissible configuration. S2 ensures that deviations from this configuration are penalized in proportion to their geometric defect. S3 ensures that the subsystem returns to saturation with superlinear speed. In FGIR, this behavior is called enaction: the passage from an unstructured dispositional state to stabilized corporeal existence. Real analyticity of the admissible protoenergy functional near the critical point guarantees the Łojasiewicz gradient inequality, converting S3 from an assumed dynamical property into a theorem for compact analytic manifolds. A stability lemma shows that the Hessian of the admissible protoenergy at a critical point is positive definite if and only if S1–S3 hold, linking the three conditions formally. Persistence implies saturation rather than the converse. Full proofs appear in Appendix B.

4. The Fundamental Law $\mathbf{E} = \mathcal{F}\mathbf{c}^2$

The SRP's conditions must now be instantiated on a concrete object. An information manifold is a smooth space whose points represent possible informational states. We use a weighted Riemannian framework (the Bakry–Émery framework [25]) as the base geometry. The manifold is a weighted space $(\mathrm{M}, g, e^{-\mathbf{I}}\,d\mu)$, where g is an ordinary Riemannian metric and $e^{-\mathbf{I}}\,d\mu$ is a volume form weighted by the information potential I. It should be noted that I denotes the qualitative information field, while italic $I(x)$ denotes its quantitative scalar evaluation at a point determined by x. The two are conceptually distinct, with the former a structured field in the four-field ontology and the latter a numerical value used in the weighted calculus.

The relevant structural consequences of this choice (the form of the weighted Laplacian, the curvature–dimension condition, and the Fisher-information structure of the weighted Poincaré and log-Sobolev inequalities) are summarized in Appendix B. The main text requires only that the framework supply the spectral-gap input needed for the stability lemma and identifies the squared-gradient exponent $\alpha = 2$ as a structural consequence of the calculus rather than a modelling choice.

A clarification bears emphasis here. The Fisher–Rao metric is not the base metric of the incorporeal manifold. A Fisher metric exists only on a named parametric family $\{\rho(\cdot \mid \theta)\}$, and the four-field ontology does not specify such a family globally. Fisher information enters solely in its proper variational role, as the carré du champ/Dirichlet form in the weighted inequalities, and via Chentsov's coarse-graining uniqueness theorem [26] for whichever parametric subfamilies arise locally.

The two structural quantities we need are the following. The global coercivity constant:

$$c := \sup_{g} \operatorname{ess\,sup}_{x} \| \nabla_g I(x) \|_g$$

measures the maximum rigidity the operative field can impose. The admissible protoenergy is the infimum, over admissible metrics, of the integral of I times the squared gradient of I:

$$E := \inf_{g \in G_{\text{adm}}} \int_M I(x) \; \| \nabla_g I(x) \|_g^2 \; d\mu_g(x). \tag{2}$$

and the protomass is:

$$\mathcal{F} := \int_M I \, d\mu_{g^*},$$

the integral of the information potential over M with respect to the energy-minimizing metric g^*. The protomass is a geometric invariant of the weighted structure, extensive under disjoint union of saturated subsystems and invariant under admissible metric deformations. The prefix proto- marks its pre-corporeal status: $\mathcal{F}$ carries no units of mass, no Lorentz-invariance property, no inertial role, until the freezing projection π identifies it with the corporeal rest mass $\hat{m}$ under the bridge conditions of Section 6 (see the companion derivation of mass genesis from photon bifurcation [27]). The squared gradient in Equation (2) is not an arbitrary exponent: the exponent 2 is uniquely forced by the joint action of ellipticity, self-adjointness, and the maximum principle [12].

In plain terms, Equation (2) asks: across every geometrically legitimate way of measuring distances and gradients on the information manifold, what is the smallest possible value of information times squared rate-of-change of information, integrated over the whole manifold? Taking the infimum, rather than fixing one metric in advance, reflects the fact that the manifold's own geometry is not given in advance either. It is itself part of what the system's persistence has to settle. A system that persists is, in effect, a system that has found the cheapest geometric configuration compatible with holding its information potential in place; systems that have not yet settled into such a configuration are precisely the proto-autopoietic or far-from-equilibrium systems flagged throughout this paper. The squared-gradient term is doing the work that rigidity intuitively suggests: a flat information potential (small ∇I everywhere) costs little to maintain, while a sharply peaked one (large ∇I) is expensive to hold in place, in direct analogy with the physical intuition that steep potential gradients require strong restoring forces.

The derivation proceeds in three stages, stated here in logical skeleton. Full details can be found in Appendix B. An upper-bound lemma bounds the admissible protoenergy above, using the fact that the pointwise gradient of I cannot exceed c anywhere on the manifold. A lower-bound lemma bounds it below by isolating the contribution from a region where the potential is approximately uniform. A Saturation Assumption then identifies a specific saturation set A* where both gradient saturation and potential uniformity hold simultaneously, collapsing the two bounds to equality. The two bounds and the saturation identification together give the Fundamental Law:

Fundamental Law: Under the Saturation Assumption (discharged for compact analytic families by the Analytic Saturation proposition), the admissible protoenergy factorizes as:

$$E = \mathcal{F}\, c^2. \tag{3}$$

The Saturation Assumption is non-trivial, and its plausibility varies by domain. For physical systems at R(6) closure, it is well-motivated. Admissible metrics extremize the Fisher information functional, and saturation corresponds to the classical ground state. For the non-physical domains catalogued in Section 8 (biological, cognitive, organizational, social), it should instead be read as identifying the regime of enactive equilibrium in the sense of [2], a modelling idealization rather than a derived consequence; whether real systems approximately satisfy it is an empirical question, addressed by the protocol of [13,28]. The law is a structural theorem for systems that have reached saturation. For proto-autopoietic or far-from-equilibrium systems, it provides a limiting benchmark rather than an exact description that characterizes only the endpoint of the gradient flow.

Describing saturation as an idealization in these domains invites a misreading that should be forestalled, since it suggests the law simply fails where the idealization is imperfect. It does not, and the reason changes what has to be assumed. The upper-bound lemma requires only that the coercivity constant be finite; it does not require saturation at all. The inequality $\mathrm{E} \leq \mathrm{Ipc}^2$ therefore holds unconditionally wherever c is finite, and saturation is precisely the condition under which that inequality becomes an equality. The coercivity law is thus the extremal case of a bound that always holds, rather than a separate claim that lapses when saturation lapses. Where saturation is imperfect the framework accordingly defines the admissibility gap, $\Delta := \mathrm{Ipc}^2 - \mathrm{E} \geq 0$, which measures the distance from critical admissibility; the degree of idealization is then a quantity of the theory rather than an unstated concession. The gap is dynamical as well as definitional. Along an admissibility flow in which the coercivity constant grows superlinearly while the enactment energy remains bounded, Δ collapses to zero in finite time, which is the rigidity-locking behavior the Łojasiewicz condition governs. Saturation is in this sense an attractor of the flow under stated conditions rather than a post imposed upon it. The flow as stated is deterministic and should be read as the mean-field form of a probabilistic dynamics, all four fields being probabilistic in this framework; the stochastic statement recovers the deterministic one in the zero-noise limit, and the qualification matters because the collapse time is then a mean rather than a guaranteed horizon for any single realization.

A probabilistic reformulation recovers the deterministic law as the zero-noise limit of a stochastic enactment layer (an Itô [29,30] SDE system for E, I_p, c^2). This is mathematically well-posed; calibration of the noise intensities remains open and is part of the empirical program of Section 10.

The Saturation Assumption is non-trivial, and its validity depends on the domain in which it is applied. In physical systems that reach R(6) closure, saturation is well-motivated: admissible metrics extremize the Fisher-information functional, and the saturation set corresponds to the classical ground state. In compact analytic statistical families, the Analytic Saturation proposition establishes the assumption rigorously.

For the non-physical domains listed in Section 8 (biological, cognitive, organizational, and social), the Saturation Assumption should instead be understood as identifying the regime of enactive equilibrium. In these settings, saturation is a modelling idealization: it holds approximately when a system has settled into a stable attractor, and it fails for systems still undergoing transition. The proto-autopoietic category is exactly the case where the assumption does not yet hold, because the system has not accumulated enough rigidity on its operative manifold to reach the saturated state.

For such systems, the coercivity law serves as a limiting benchmark (the endpoint of the gradient flow) rather than an exact description of the flow itself. Practically, this means the law is exact for systems at or near saturation, approximate for systems approaching it, and diagnostic (a target rather than a description) for systems far from it.

Whether a given system approximately satisfies the assumption is an empirical question. This is addressed by the critical-slowing-down protocol of [28], which detects an approach to a stability boundary through rising relaxation times, increasing autocorrelation, and delayed recovery, and by the resonance-signature protocol of [13].

The Fundamental Law is a genuine theorem and non-circular by construction. Enactment E is defined directly from (M, I, g_{adm}), and the factorized form $E = \mathcal{F}c^2$ emerges only from the Saturation Assumption, not from a definitional shortcut. It is consistent with the information-geometric program of [31] and Frieden and Soffer's variational EPI foundations with its metacybernetic application to autopoiesis [32,33], set out systematically in Frieden's treatment of extreme physical information [34] and extended to biological persistence by Yolles and Frieden's study of viral living systems [11]. The ratio J/I of Section 2 is the reciprocal of the autopoietic efficacy ratio I/J of [33], and the sustentative coupling K↔(J–I) is the information change of the same treatment, so that two of the three coupling channels are EPI relations restated geometrically. The correspondence is definitional: no estimator of I/J is used in this paper, and no claim is made that any operational proxy realizes the relation. Three summary verdicts, restated at the end of the paper, consolidate the paper's epistemic status: the Fundamental Law itself is a theorem; the bridge to physics is conditional (T-BRIDGE depends on adopted bridge conditions); and the four-field ontology is a working commitment grounded in the framework's axiomatic basis (A1–A10), not a derived theorem. These are stated here so a reader can locate, in one place, exactly what is and is not established.

This answers a charge of circularity, but not a further question that lies beneath it, and the distinction is worth marking. Showing that E is defined independently of its factorized form establishes that the derivation does not assume what it proves. It does not establish that a system satisfying S1 to S3 is thereby autopoietic in the sense biologists and organizational theorists use the term. That further claim requires an account of how gradient saturation on the information manifold corresponds to the maintenance of a boundary by the processes that the boundary in turn makes possible. The account runs through the process intelligences and their coupling operators, and we give it here rather than in the conclusion so that the reader carries it through the sections that follow.

In the metacybernetic architecture [1] each recursion order is joined to the one beneath it by an integrity couplet, enacted by a process intelligence: the first couples R(1) to R(2) and is autopoiesis; the second couples R(2) to R(3) and is autogenesis, which stands to the autopoietic couple as autopoiesis stands to the operative field itself. The efficacy of a couplet is the capacity of the higher order to constitute, shape, and constrain the informational dynamics of the lower, and that efficacy is not a quantity separate from the geometry. Curvature is the rate of change of the information potential across the manifold, and at each order the curvature is set by the couplet: strong coupling produces high curvature, weak coupling flat geometry and poor constitutive capacity. The gradient of the information potential is therefore the constitutive efficacy of the process intelligence, and not an analogue of it.

Three correspondences follow, and together they answer the question. Gradient saturation, in which the gradient norm equals c almost everywhere on the saturation set, is the condition under which the autopoietic couple operates at the greatest constitutive efficacy its admissible geometry permits, uniformly across the active set; this is the geometric form of self-production sustained at capacity. Defect rigidity, the quadratic cost that deviations carry, is boundary maintenance: a boundary is a locus at which deviation is penalized, and a system whose deviations cost nothing has no boundary to maintain. The quadratic penalty is what holds the operative organization distinct from its surroundings, which is the work that the network of processes is required to do on Maturana and Varela's account. That this reading is not imposed is confirmed by the place defect rigidity occupies in the derivation, being the interior analogue of the quadratic boundary penalty, the one functional expressing at the boundary what the other expresses throughout the interior. A condition already understood as the interior counterpart of a boundary functional is not being pressed into a boundary reading; it is being read as what it is. Rapid locking, the superlinear return governed by the Łojasiewicz inequality, is the boundary's re-production rather than its mere persistence; autopoiesis is a continuous production of the boundary by the processes the boundary makes possible, and superlinear return is the geometric expression of that continuity, distinguishing a living boundary from an inert one. The recursion then completes the account, since the second couplet stands to the first as the first stands to the operative field, so that boundary maintenance at each order is the same structure applied to the couple below it. Each process intelligence is legible through the Fisher functional, high information marking sharp discrimination and efficacious constitution and low information marking flatness, so the coupling efficacies are Fisher-filtered throughout. One boundary should be marked, and marking it requires separating two questions. Whether autopoiesis is restricted by definition to molecular biology: Maturana and Varela held that it was, and Maturana later reaffirmed the restriction in qualified terms [35], but that limitation has been overcome in the sociocybernetic literature [36] and in Razeto-Barry's reformulation [37], so that autopoiesis applies to living systems whether biological, chemical, psychological, mechanical or social, which is the generic reading adopted here [33]. Whether the generic statement licenses application to any system whatever: it does not, and the governing condition is not a matter of judgement. In [33] it is that the parametric representation of context be bounded in error and that autopoietic efficacy be high; the present framework states the same requirement metrically, as saturation on a well-defined constraint class (S1–S3), R(6) closure for the corporeal reading, and the tracking conditions of [13]. The correspondence is therefore not an analogy drawn to a biological original but the information-geometric expression of the same closure, carrying its conditions of application with it.

5. Fisher-Generative Specialization: $E = I_{\mathrm{P}}c^2$

The Fundamental Law holds for any information manifold satisfying the Saturation Assumption. Obtaining the FGIR-specific coercivity law $E = I_{\mathrm{P}}c^2$ requires specifying which weight is in view. Where a statistical subfamily $\{\rho(\cdot \mid \theta) \colon \theta \in \Theta\}$ is in fact parametrized, the Fisher–Rao metric is, by Chentsov's theorem [26,38,39], the unique coarse-graining metric (up to a positive scalar) invariant under all Markov morphisms, the natural variational structure on that subfamily. Its connection to the SRP's defect-rigidity condition S2 is mediated by the Kullback–Leibler divergence, whose local quadratic expansion is exactly the squared-defect structure S2 demands. The base geometry of the incorporeal manifold, however, remains the Bakry–Émery weighted metric of Section 4; the Fisher–Rao metric enters only as the coarse-graining variational metric on whichever parametrized subfamilies are locally in view (for example, the qubit Bloch-sphere test of [40]). Full mathematical detail is provided in Appendix B.

Identifying the abstract information potential I of Section 4 with the FGIR-specific integrated informational potential I_p has two components. Mathematically, the Fisher information functional on Θ is an instance of the abstract scalar field I satisfying the standing assumptions of Section 4. Ontologically, Ref. [12] established that within FGIR the operative I-field carries a Fisher Information Topology and performed the embedding $I \to I_p$, distinguishing the general functional from the program-specific density parameter. This embedding allows the invariant $E = I_p c^2$ to be read simultaneously as a geometric identity on the incorporeal manifold and as the precursor to its corporeal interpretations upon freezing. The identification $\mathcal{F} = I_p$ is structural rather than notational: $\mathcal{F}$ is defined on the incorporeal manifold M, while I_p is defined on the operative manifold M_I. The embedding of Section 2 identifies the operative stratum of M with M_I and preserves the weighted measure, so both integrals compute the same invariant viewed from two ontological registers.

The global coercivity constant $c = \sup \|\nabla_g I_p\|_g$ measures the maximum rigidity the operative field can impose. It is an intrinsic quantity, the supremum of the gradient norm of the information potential over the admissible class, so its later identification with a propagation speed interprets a quantity already present rather than importing one. Since I_p satisfies all the type and regularity conditions imposed on I in Section 4, the Fundamental Law applies without modification, giving the specialized coercivity law directly:

$$E = I_p \, c^2. \tag{4}$$

This invariant is the structural law on which the framework's applied layers depend, including the Constitutional Persistence Cycle, the Wuxing reconstruction, the Beer VSM application, and the fractal-seed construction [1]. The present derivation supplies the proof that underwrites those dependencies, replacing stipulation with a formally established result.

6. The Freezing Projection: $E = I_p c^2 \;\Leftrightarrow\; \hat{E} = \hat{m}\hat{c}^2$

The coercivity law has been derived on the incorporeal information manifold; Einstein's law belongs to four-dimensional Lorentzian spacetime. We formalize the freezing projection π as a topological phase transition of the Bakry–Émery weighted metric. In the incorporeal regime, the weighted metric is fluid, multiple parametric trajectories coexist, curvature is distributed and non-localized, and the manifold admits superposition. When the system loses K-governance, that distributed curvature collapses into rigid, localized Riemannian curvature concentrated at a point, the metric analogue of a liquid-to-solid transition. The corporeal quantities $\hat{m}$ and $\hat{c}$ are geometric residues of this collapse, not arbitrary relabelings, which is what distinguishes this account from a notational trick. This section states the projection and its three bridge conditions; the deep differential geometry (Picard–Lefschetz monodromy, Berezin–Pfaffian entity-collapse, the full CPC construction) is deferred to Appendix B.4.

An intuitive perspective may be useful before introducing the formal machinery. Think of the incorporeal information manifold as a superpositional option space, a probability cloud of co-possible configurations that have not yet precipitated into actuality. In FGIR, superposition is contextual rather than amplitude-based: multiple admissible configurations remain simultaneously viable because the K-field itself is in contextual superposition, sustaining several co-possible sustentative boundary-condition frameworks at once. This higher-order superposition keeps the option-space aloft. Weighted-metric fluidity is the formal expression of this contextual suspension: curvature is distributed rather than localized, the manifold remains non-rigid, and no single trajectory is selected as the configuration. Freezing is the moment when that governance fails or saturates at a locus, and the suspended option-space rains out into a single, rigid, metric-bearing geometry, a nimbus collapse of dispositional curvature. The residue of that precipitation (definite curvature concentrated at a point rather than diffuse curvature spread across the manifold) is what corporeal physics calls mass. This is why $\hat{m}$ is a residue of collapse rather than a relabelling of I_p, where the freezing projection performs the geometric work of converting distributed indeterminacy into a localized, metric-bearing fact, and the resulting mass is the leftover rigidity of that conversion, not merely the same number expressed in different units.

It is useful to separate π conceptually into two stages, even though it is defined as a single operation. Enactment E, defined in terms of I_p, and c, are derived from the incorporeal manifold under the SRP. When they are forced into a persistent, boundary-satisfying configuration under the Saturation Assumption, they retain the relation:

$$E = I_\mathrm{p}\, c^2.$$

Under the bridge conditions stated below, the enacted quantities become embedded in C_impl, thereby acquiring the conditioning appropriate to C_expl. The three bridge conditions are formulated with respect to the freezing map, which defines the nature of freezing itself. A freezing map specifies how an informational quantity on the incorporeal manifold is reassigned to its corporeal image under the projection:

$$\pi : \big(M, I_p, c\big) \longrightarrow (\hat{m}, \hat{c}),$$

a topological phase-transition of the Bakry–Émery weighted metric, where phase denotes a dispositional geometric regime in the FGIR sense. This geometric collapse satisfies BC1.

BC1 (Dimensional-reduction).

π reduces the degrees of freedom to a four-dimensional Lorentzian manifold, with the collapse of the fluid weighted metric into localized curvature at a point.

BC2 (Constitutional-identification/Tensor–Rigidity Correspondence).

$$\pi\big(I_p\big) = \hat{m}.$$

Constitutional persistence condenses into rest-mass, the geometric residue of the metric-collapse; the integrated information-weight yields an extensive, conserved scalar under freezing.

BC3 (Energy-preservation/ι-convergence).

$$\pi(E) = \hat{E}, \pi(c) = \hat{c}.$$

Elsewhere in the FGIR framework, ι denotes informational inertia, defined as the system's resistance to directional change on the weighted informational manifold. It is the rigidity variable that governs how rapidly or slowly perturbations propagate through the J–I–K triad. Under the specific condition of enactment into corporeality, when the freezing projection π is applied, and constitutional closure at $\mathrm{R}(6)$ is active, this same scalar acquires a second role: the rigidity that governs local perturbation–propagation also governs the rate at which the weighted metric contracts under recursive closure. In this context, ι functions as the recursion index. BC3 does not redefine ι, but imposes a convergence requirement on informational inertia in its constitutional role. The recursion index ι must stabilize (by reaching a fixed point) so that the invariant

$$E = I_p c^2$$

is preserved across the freezing projection. Energy-preservation and ι-convergence are therefore two expressions of a single structural condition required for the T-BRIDGE equivalence: only when informational inertia converges can the corporeal coherence-constant $\hat{c}$ be well-defined.

The numerical value of the enaction-integral is preserved, its ontological type changing from generic to metric-bearing upon residence in C_{impl}. The coherence-conductance freezes as the invariant signal-speed, with

$$\hat{c} = c$$

within the corporeal field. Preservation of E and c is governed by ι-convergence, the closure condition ensuring that informational inertia ι stabilizes under the transition from the weighted metric to its corporeal image.

BC1–BC3 are adopted as working conditions at the closure level R(6), not proven within this paper; their formal discharge from R(6) recursive-closure is the central open problem of the framework. The three conditional discharge theorems (Lorentzian-signature, time-translation-symmetry, and characteristic cone uniqueness) jointly constrain the corporeal image of the weighted metric. A Lorentzian signature ensures that the collapsed geometry carries exactly one timelike direction. Time-translation-symmetry guarantees that the enacted invariant admits a conserved energy interpretation. And characteristic-cone-uniqueness secures a single, non-branching causal cone, eliminating competing signal speeds in the corporeal field. Taken together, these conditional theorems reduce the bridge-architecture to four sharply bounded conjectures.. We will not elaborate on them here because, for the cybernetic thesis of this paper, the bridge-conditions enter only as the condition that licenses the physical reading of the invariant. The cybernetic threshold itself (the question of when a CAS becomes a CAAS) does not depend on their discharge.

7. T-BRIDGE: Conditional Freezing Equivalence Theorem

Working within the FGIR fractal seed framework (the four-field ontology, the coercivity law $E = I_{\mathrm{p}}c^2$ at the incorporeal level, and the freezing projection π acting at R(6) closure on C_{impl}), the T-BRIDGE theorem formalizes the equivalence sketched in Section 6. We shall state it here for completeness, but readers whose interest is cybernetic rather than physical may proceed directly to Section 8.

Proposition 2. *On C_{impl}, under π satisfying BC1–BC3, the informational coercivity law and the relativistic mass–energy law are the same equation in different ontological representations:*

$$E = I_p\, c^2 \iff \hat{E} = \hat{m}\, \hat{c}^2. \tag{5}$$

Proof. Start from the incorporeal law $E = I_p c^2$. Apply π and use BC3 (energy preservation):

$$\hat{E} = \pi(E) = \pi(I_p c^2).$$

By the functoriality of π on products and BC2:

$$\pi(I_p c^2) = \pi(I_p) \cdot \pi(c)^2 = \hat{m}\, \hat{c}^2$$

(using $\hat{c} = c$ within C_{impl}, by BC3). Hence $\hat{E} = \hat{m}\hat{c}^2$. The converse follows by running the same steps in reverse from $\hat{E} = \hat{m}\hat{c}^2$ on C_{impl}. BC1 fixes the dimension and signature of the corporeal target; BC2 and BC3 identify the corporeal images of I_{p} and c. □

In other words, the T-BRIDGE establishes that if the freezing projection π satisfies the three bridge conditions BC1–BC3, then the informational coercivity law $E = I_{\mathrm{p}}c^2$ and the relativistic mass–energy law $\hat{E} = \hat{m}\,\hat{c}^2$ are the same equation in two distinct ontological registers. The equivalence requires the non-trivial identifications $I_\mathrm{p} \to \hat{m}$ and $c \to \hat{c}$, grounded respectively in the Tensor–Rigidity Correspondence (BC2) and the iota-convergence theorem (BC3). The distinction between what T-BRIDGE establishes versus what it assumes matters for the cybernetic contribution of this paper. T-BRIDGE establishes an equivalence, a statement of the form "if BC1–BC3 hold, then the two laws are the same equation." What it assumes is the bridge conditions themselves, which are adopted as working conditions at R(6) closure, not proven here. The cybernetic threshold $E = I_\mathrm{p}c^2$ (the question of when a CAS becomes a CAAS) is established independently of the bridge (Sections 3–5) and does not depend on its discharge. The bridge simply records that the same invariant can be read as physical mass-energy when the bridge conditions hold. Readers whose interest is cybernetic rather than physical may therefore treat Sections 6 and 7 as ontological grounding for the framework, not as a load-bearing part of the paper's central contribution.

8. Corporeal Projections and the Operationalization of Cost

$E = I_\mathrm{p}c^2$ is a general law of enaction, holding on any information manifold satisfying the SRP regardless of domain. Its corporeal projections fall into two structurally distinct categories. Direct freezing projections collapse $\left(E, I_\mathrm{p}, c\right)$ onto a given corporeal field by direct application of π. The Constitutional Persistence Cycle (CPC) derivations present equations reached by the full recursive dynamics: a rigidity lock, an entity-collapse step, and a monodromy reset. The CPC–Telegrapher Theorem, which establishes that these three steps generate the telegrapher's equation as the common ancestor of the Schrödinger, diffusion, and Dirac equations, is a complete derivation ($\mathrm{T^E}_0$) conditional on the entity-collapse step; the explicit Berezin–Pfaffian construction of that step is itself complete only for block-diagonal/normal operators ($\mathrm{T^E}_0$) and remains conjectural in the general case ($\mathrm{T^E}_2$; Appendix B.4). The Nelson stochastic-mechanics identification that links the CPC output to quantum mechanics remains an externally borrowed result ($\mathrm{T^E}_1$). Table 2 catalogues the direct-projection rows.

Table 2. Direct freezing projections of the coercivity law, reclassified into four epistemic tiers.

Corporeal Context	I_p Maps to	c Maps to	Resulting Law	Tier
Special relativity (flat 4D)	Rest mass $\hat{m}$	Speed of light c	$\hat{E} = \hat{m}\hat{c}^2$	$\mathrm{T^E}_1$ (conditional on BC1–BC3)
General relativity (curved)	Local mass density	Local signal speed ($g\mu\nu$-dependent)	$\hat{E} = \hat{m}\hat{c}^2$ locally	$\mathrm{T^E}_1$ (conditional on BC1–BC3)
Quantum mechanics (general)	Quantum Fisher information density	Coherence propagation speed ($\hbar$/m)	$E = FT{\cdot}\hbar^2/m^2$; Cramér–Rao bound	$\mathrm{T^E}_1$ (conditional; CPC-Telegrapher $\mathrm{T^E}_0$, block-diagonal case)
Thermodynamics	Fisher information of thermal distribution	Thermal velocity $\sqrt{(kT)}$	$E = kT{\cdot}I_\mathrm{p}$	$\mathrm{T^E}_2$ (CPC overdamped limit)
Condensed matter	Quasiparticle density of states	Phonon/Fermi velocity vF	$E = I_\mathrm{p}{\cdot}vF^2$	Heuristic

Table 2. *Cont.*

Corporeal Context	I_p Maps to	c Maps to	Resulting Law	Tier
Bose–Einstein condensate	Coherence fraction	Bogoliubov sound speed	Ground-state energy	Heuristic
Black-hole thermodynamics	Bekenstein entropy density	Hawking temperature/Planck scale	$S = A/4l_p^2$	T^E_1 (consistent with established black-hole thermodynamics; Enactment paper Section 4 worked test)
Biological/autopoietic	Organismal Fisher information (regulatory precision)	Neural/chemical signal speed	Metabolic cost of enacted existence	T^E_2 (derived via companion Enactment paper πbio, Section 7; testable per Section 3 protocol)
Cognitive/enactive	Cognitive Fisher information	Neural propagation speed	Neural energy cost of coherence	T^E_2 (derived via companion Enactment paper πcog, Section 8; testable per Section 3 protocol)
Organizational (Beer VSM)	Requisite variety (Ashby)	Information transmission rate	Organizational/management cost	T^E_2 (candidate indicators, heuristic per Section 8; testable via the resonance-signature protocol of [13])
Social/political	Social coherence density	Cultural propagation speed	Cost of maintaining social order	T^E_2 (derived via companion Enactment paper πsoc, Section 9; testable per Section 3 protocol)
Cosmological/inflationary	Vacuum fluctuation density	Inflationary propagation speed	Inflationary energy density	Heuristic

Table 2 shows that the twelve rows are not epistemically equivalent, and the four-tier reclassification makes this explicit: T^E_1 rows are conditional theorems (pending discharge of the bridge-condition) or formal consistency checks; T^E_2 rows are rigorously derived CPC limits or derived projections with a specified, testable proxy pathway; and only the remaining three rows, lacking any such derivation or pathway, are heuristic, speculative formal isomorphisms illustrating the framework's range and not empirical confirmations. The relativity rows are conditional theorems (T^E_1), pending discharge of the bridge conjectures of Section 6. The quantum-mechanics row is a conditional theorem (T^E_1): the CPC-Telegrapher Theorem from which it derives is a complete derivation (T^E_0) conditional on the entity-collapse step, whose explicit construction is complete for block-diagonal/normal operators (T^E_0) but conjectural in the general case (T^E_2; Appendix B.4); in addition, the Nelson stochastic-mechanics identification that links it to quantum mechanics remains externally borrowed. The thermodynamics row is a rigorously derived overdamped limit of the CPC telegrapher equation (T^E_2), consistent with Landauer's principle linking informational and thermodynamic cost [41], and its black-hole sub-case is separately consistent with established black-hole thermodynamics, reproducing the Bekenstein–Hawking relation exactly (T^E_1, Section 4 worked test). The biological, cognitive, and social rows are T^E_2: derived via the companion Enactment paper's cross-disciplinary domain applications [42] (π_{bio}, π_{cog}, and π_{soc} projections, Sections 7–9) and testable in principle via the Section 3 protocol, though not yet empirically confirmed. The organizational (Beer VSM) row joins this T^E_2 group on the same basis: [13] states the four-condition resonance-signature protocol against which candidate C-field observables are testable, the organizational indicators offered below being proposed here rather than carried over from it, and it carries a qualitative empirical precedent, Rautakivi and Yolles' ASEAN diagnosis, which performs the same freezing move against real institutional data (see Empirical program, below). None of this is yet empirical confirmation, but it is a specified, testable pathway rather than a bare formal isomorphism. Only the remaining three rows (condensed matter, Bose–Einstein condensate, cosmological) are properly heuristic, speculative formal isomorphisms illustrating the framework's range of application without empirical confirmation or a specified proxy pathway. All heuristic and $T^E{}_2$ rows in non-physical domains inherit an unresolved dimensional-identification problem, to which we now turn.

A social systems theorist reading Table 2 will reasonably ask how I_p and c might be measured in an organization? If organizational energy or the cost of maintaining social order has no measurable units, the theory risks being dismissed as metaphorical. The framework's response is the Proxy Tracking Theorem of [13], which we retain as the bridge from the formal apparatus to operational practice. This shows that C-field observables (the corporeal quantities an analyst can actually measure) track incorporeal field dynamics through a governed sufficient-statistic cascade rather than mere correlation. The theorem is valid whenever four jointly necessary conditions hold: (i) a discrete, finite dispositional configuration space; (ii) absolute continuity of the operative manifold with respect to that space; (iii) a resonance condition on the H-manifold's resonant subspace; and (iv) compactness of the operative manifold over the observation window. When these conditions hold, a domain's $I_\mathrm{p}c^2$ law is validated not by locating an antecedent physical unit system but by detecting a resonance signature in C-field time-series data. This addresses the dimensional-identification problem with a checkable empirical protocol, within the limits set out next.

In practical terms, a systems theorist can proceed as follows. The integrated informational potential Ip is a measure of how much regulatory information the operative field is holding in place. In an organization, candidate C-field indicators include communication-network density (the richness of the operational communication graph that carries the regulatory signal), the diversity of decision rules actively in force, and the precision with which deviations from operating norms are detected. The coherence conductance c is a measure of how quickly informational structure locks into place; candidate indicators include decision-making latency (how rapidly a perturbation to operating norms is registered and acted upon), the speed at which a policy change propagates through the organization, and the inverse of the mean response time of the regulatory layer. These indicators are not arbitrary: each is a candidate observable of the incorporeal quantity the framework identifies, and each is testable against the resonance signature the Proxy Tracking Theorem predicts. Their status, however, is constrained in a way the framework itself specifies, and we state it next.

The constraint is the third bridge condition, and it is stronger than a shortage of data. As defined in Section 4, c is a supremum of a gradient norm on the information manifold, carrying units of information per unit informational arc length. Decision latency and propagation speed are rates, in units of information per second. The passage between the two is not a calibration exercise; it is precisely what BC3 asserts, namely that under R(6) closure the supremal gradient converges to a finite signal speed by convergence of the informational inertia. BC3 is not discharged. Further, the convergence it names is a property of R(6) closure, and the organizational domain does not exhibit that closure, so the informational inertia there has no fixed point. It remains defined, and a generalized uncertainty relation continues to hold with it, but nothing in the structure holds it at a value. The conversion between c and any measured rate therefore has no value that persists: a calibration obtained at one time carries no warrant at another, because no attractor returns the system to it. Since quantitative measurement requires a conversion constant that survives, the quantities named above are heuristic indicators of enactment cost and not measurements of c or of Ip. We say so plainly, because the alternative reading, on which better instrumentation would close the gap, misdescribes what the framework claims.

The constraint bears on absolute magnitudes, and it is worth being equally exact about what escapes it, since the framework does supply an organizational quantity that survives. The admissibility gap of Section 4 may be normalized, giving a saturation deficit Δ/Ipc^2 lying between zero and one. Being a ratio of like quantities it is dimensionless, and a dimensionless ratio requires no unit identification: it is invariant under whatever conversion BC3 would eventually supply, and is therefore well defined in domains where that conversion is not. What can be assessed organizationally is accordingly not how large Ip or c may be, but how far a system stands from critical admissibility, and whether that distance is widening or closing. The dynamical result quoted in Section 4 gives this an observable signature, since approach to saturation is accompanied by the divergence of relaxation time that the critical-slowing-down diagnostics of the companion work are designed to detect [28], and those diagnostics are themselves ratio-based rather than absolute. This is the operational scheme the framework can presently support for organizations, and we prefer to offer it than to offer absolute proxies that the preceding paragraph disallows.

One consequence deserves emphasis, because it explains why the deficit rather than the constant is the tractable quantity. The coercivity constant is a supremum taken over the class of admissible metrics, and that class is not fixed independently of circumstance: the operative substrate constrains which metrics are admissible, so that different operative and corporeal configurations impose different constraint classes and yield different suprema. The global admissible class is itself an idealization, the operative construct being a family of local constraints. Corporeal context therefore enters the determination of c, rather than merely furnishing occasions on which a pre-existing c is displayed. This is why no single organizational value should be sought for it, and why the absence of a fixed point for the informational inertia in domains short of R(6) closure is not a further difficulty but the same one differently expressed: where context continually re-selects the constraint class, nothing settles the supremum. The deficit is stable against this because it is a ratio formed within whatever constraint class currently obtains.

This also explains why a perturbation cannot be treated as disturbing one condition while leaving the others fixed, and the route it takes is worth tracing because it is not the one a reader might assume. A change of context is a perturbation of the corporeal explicate order, and it reaches the informational fields by a route with two distinct stages, neither of which is a push from outside. Within the corporeal field, the explicate stratum re-enfolds as updated parameters of its own implicate stratum. This is the lateral register, which obtains at the B-manifold gate of each field independently, there being one such gate for the dispositional, operative, sustentative, and corporeal fields severally. The corporeal field is then coupled to the incorporeal triad across realms at the H-manifold, which is the transverse register. The two registers are orthogonal and must not be conflated: the lateral couples implicate to explicate within a field, the transverse couples realms. What obtains at either gate is informational entanglement, the term being used advisedly and in its own right: a structural interaction governed by Fisher-informational geometry, consisting in mutual constraint and co-determination among statistical-geometric objects on a Riemannian information manifold under the Fisher–Rao metric. The manifold is the gate; the entanglement is the correlational phenomenon occurring at it. What passes between the corporeal and the incorporeal is therefore correlation among informational structures, and not transmission of anything between them. The dispositional, operative, and sustentative fields are perturbed together in consequence, with higher recursion orders affected to a lesser degree. The conditions of the Saturation-Rigidity Principle accordingly move together rather than singly, which is why the framework treats them as simultaneous rather than as a checklist, and why corporeal context is a determinant within the geometry rather than an external circumstance acting upon it.

The relation of this to quantum entanglement should be stated with care, since the two are neither the same thing nor merely alike. Informational entanglement is not quantum entanglement in the Hilbert-space sense: there is no joint wavefunction, no non-separable state and no amplitude structure linking the incorporeal fields, and the recursive coupling among them must not be read as quantum entanglement in any register. Where the coupled objects are literal quantum states, the informational relation is required to specialize to quantum entanglement rather than merely to resemble it, since the Fisher–Rao metric restricted to a quantum statistical manifold is the quantum Fisher information, which is an established witness for genuine multipartite entanglement [43]. Whether the corporeal-field construction used here specializes correctly to that quantity in the quantum case, and so inherits the witnessing property, is not derived in this paper and remains open. We record the requirement and its undischarged status together, because the first without the second would claim a reduction we have not performed.

How the perturbation then propagates among those fields is equally specific, and it bears on the tracking argument the paper relies upon elsewhere. Propagation is transitive through the process intelligences: each intelligence enacts a map between fields, and a perturbation traverses the architecture as a composition of those maps rather than by direct field-to-field action. This is not merely a picture of the transmission but the mechanism on which the Proxy Tracking Theorem turns, since the resolution map characterizing the dispositional-to-operative transition and the map across the transverse gate are each sufficient statistics for the parameters at issue, and it is by the transitivity of sufficient statistics that their composition carries the tracking relation intact from the dispositional simplex to the corporeal manifold [13]. The same transitivity that licenses corporeal diagnostics to track incorporeal structure licenses corporeal perturbation to reach it. Tracking and perturbation are the one relation read in opposite directions, which is why a framework that admits the first cannot treat the second as external.

The consequential step is what follows from the dispositional perturbation in particular, since this is what makes the admissible class move rather than merely tremble. Perturbation of the dispositional field is the occasion for new Fisher information topologies, the structured actualizations of operative potential whose emergence the sustentative field orchestrates through the process intelligences, transforming dispositional perturbation into stable thematic elements [12]. New topologies are new operative options. They enlarge or restrict what the operative field can hold in place, and since the admissible geometry class is precisely the class of metrics compatible with what the operative substrate can sustain, a change in the available topologies is a change in that class, hence in the supremum taken over it, hence in the coherence conductance itself. The path from a change of circumstance to a change in c therefore runs through the dispositional field and the topologies it makes available, and not by any direct action of circumstance upon the operative geometry. This is the sense in which the framework treats context as generative rather than merely as a boundary condition, and it is why we hold that no single organizational value of c should be sought: the quantity is constituted anew as the topologies change.

This is a foundational commitment of the framework rather than a reading imposed here. Enactment is held to be driven by context, that is by boundary conditions and operative geometry, and not by the Fisher–Rao metric, which is an instrument for computing admissibility rather than a cause of anything. The commitment carries a notational con-

https://doi.org/10.3390/systems14091132

sequence observed throughout: the imperative ratio of dispositional to operative field is a qualitative relation between incorporeal fields and not a computable equation, so quantitative work proceeds through scalar proxies marked as such and not to be confused with the fields they estimate. The concern that variables change meaning under cross-domain mapping is therefore met by a principle the framework already holds, and not by a concession made under pressure.

It is important to be clear about what these proxies are and are not. They are not empirical confirmations that organizational cohesion literally obeys an energy-conservation law; they are operational handles on a structural isomorphism. The heuristic projections in Table 2 are formal isomorphisms, not empirical confirmations, and should be read as a demonstration that the same mathematical form recurs across domains that FGIR treats as coupled by a common informational substrate (a structural claim worth taking seriously in its own right) without over-reading it as a claim that, say, an institution's resilience is literally measured in joules. What the proxy approach does provide is a checkable protocol: a domain is added to the validated set not by stipulation but by carrying out the four-condition resonance test on its C-field data.

To date, only one domain has been carried through this protocol in full: Yolles' Iran–Israel/US conflict study [16] was the framework's first structured application, with three of the four conditions satisfied and the fourth pending prospective verification. That study is an unrefereed preprint by one of the present authors, and it exhibits the predicted pattern without measuring the quantities; it therefore evidences the existence of a freezing-consistent diagnostic methodology rather than the law itself. The full four-condition protocol, the resonance-signature test methodology, and the calibration apparatus for C-field proxies are set out in [13]. Readers seeking to apply the proxy approach to a new domain should consult that paper for the operational details, which are beyond the scope of the present derivation. A minimal falsifiable footprint is established for critical-slowing-down predictions in non-compact, long-memory domains, providing applicability without prior calibration or resolution of the bridge conditions [28]. The organizational domain has a related but distinct precedent: Rautakivi and Yolles' two-part diagnosis of ASEAN [44,45] performs, in practice, exactly the freezing move Section 6 formalizes. It maps an organization's incorporeal dispositional and operative structure onto observable institutional behavior and validates the mapping against survey data and twenty years of documented institutional outcomes. It does so qualitatively, through trait diagnosis rather than through the quantitative apparatus of [13]; we cite it as evidence that a freezing-consistent diagnostic methodology for the organizational domain already exists and has been tested against a real case, one rung below full quantitative calibration rather than absent altogether.

The scope of what this establishes should not be overstated, and we prefer to state it rather than leave it to be inferred. No domain, physical or otherwise, has been subjected to a quantitative test of the coercivity threshold. The single structured application satisfies three of the four conditions of our own protocol, is a preprint by one of the present authors, and is qualitative in the sense that matters here: it exhibits the predicted pattern without measuring the quantities. The law is therefore advanced as a structural theorem within FGIR, accompanied by a specified but as yet unexecuted test program, and not as an empirically supported law of the domains it addresses. A related concern deserves a direct answer. A theory whose conditions are numerous, and whose failures can always be attributed to an unmet condition, is unfalsifiable in practice, however well-formed it may be in principle. The protection against this is that the conditions be checkable independently of the outcome they are invoked to explain, and we accept that obligation. Saturation is to be assessed by the critical-slowing-down diagnostics of the companion work and

R(3) closure by the architectural criteria of the companion metacybernetics paper, each established before the threshold test is run rather than after it. A failure of the threshold test in a system independently shown to satisfy both would count against the law, and we would regard it as counting.

The connection to Ashby's Law of Requisite Variety [46] is worth drawing out explicitly, because it shows precisely what the coercivity law adds to a foundational result in cybernetics. Requisite variety states that a regulator's variety must be at least as great as the variety of the disturbances it must absorb; it is a combinatorial condition on the space of possible states; its companion result, that every good regulator of a system must be a model of that system, states the same requirement in modelling terms [47]. The coercivity law adds a second, energetic condition that requisite-variety accounts typically leave implicit: possessing sufficient variety is not, by itself, sufficient for a regulator to persist, if the rigidity cost of holding that variety in a stabilized, enacted configuration exceeds what the regulator's operative field can sustain.

Concretely, a management system might in principle possess requisite variety (with enough procedural richness to match every disturbance its environment can present) while still failing as a going concern, because the rigidity cost of continuously re-enacting that variety under real perturbation exceeds the coercivity threshold available to it. The system has the right shape (architectural sufficiency) but cannot pay the cost (corporeal viability). This reframes an old cybernetic intuition that viable systems need enough slack as well as enough variety. In the vocabulary of the present framework, slack, on this reading, is a colloquial name for the margin between a system's actual $I_p c^2$ and the coercivity threshold the SRP requires it to clear. A system with ample variety but sub-threshold $I_p c^2$ will exhibit exactly the proto-autopoietic signature: rich operational behavior while external boundary conditions hold, collapse into incoherence when they are withdrawn.

This highlights a practical contribution of the framework to systems science. Ashby's Law tells the designer whether the regulator has enough combinatorial capacity. The coercivity law tells the designer whether the regulator has enough structural rigidity to use that capacity over time. The two together, a combinatorial sufficiency and energetic sufficiency, give a two-axis criterion for viable system design that neither supplies alone.

The status of this correspondence should be stated exactly, since as drawn so far it is an interpretive parallel. Requisite variety is a combinatorial quantity, the logarithm of a state count; Ip and c are informational-geometric quantities on a weighted manifold. The term slack, used above, is a colloquial gloss rather than a defined quantity of the framework. A formal connection can nonetheless be proposed, and we set one out below as a conjecture rather than a theorem, since it turns on an identification we do not discharge here.

Accessible variety may be defined from the Fisher structure rather than from a state count. On a parametric subfamily of the operative manifold with Fisher information matrix having eigenvalues λ_1 to λ_d, and with κ the positive constant of the quadratic boundary penalty $\Lambda(I) = (\kappa/2)I^2$, define the accessible variety as $V_{acc} = ½ \Sigma \log2(1 + \lambda_i/\kappa)$. This counts the configurations a parametric family can separate at a resolution fixed by κ, and it is the information-geometric counterpart of a state count. The notion of separation is to be understood in the sampling sense the framework requires throughout, probabilities here being frequentist: two parameter points are distinguishable when the sampling distribution of an estimator separates them, the Fisher information being the precision governing that separation, so that the count is a statement about the resolving power of the model and not about any distribution over its parameters. No prior is invoked and none is required. Chentsov's theorem [26] makes the Fisher–Rao metric the unique choice invariant under sufficient statistics, and so the unique metric under which such a count is well defined at all [31]; the regularized log-determinant is the standard form of that count. Ashby's condition then reads $V_{acc} \geq V_{req}$, where $V_{req} = \log 2\, N$ is the logarithm of the number of disturbance classes the regulator must distinguish.

The choice of κ as the scale against which the Fisher eigenvalues are measured is not arbitrary borrowing. Defect rigidity is the interior analogue of that same boundary penalty, the boundary functional expressing at the manifold boundary what defect rigidity expresses throughout the interior, so κ is already the constant fixing what counts as a costly deviation in this geometry. It is therefore the appropriate resolution at which configurations are counted as distinguishable, and no further parameter is introduced.

Two features of this formulation deserve notice. The first is that both sides are dimensionless. V_{req} is the logarithm of a count, and V_{acc} is the logarithm of a ratio of like quantities, the eigenvalues of the Fisher matrix being measured against the boundary-penalty constant of the same structure. The comparison therefore requires no unit identification and is not obstructed by the constraint of Section 8, which bears on absolute magnitudes alone. The second is that the condition so expressed is a condition on the rank and volume of the operative geometry, that is, on how many configurations the system can distinguish, and is silent on whether those distinctions can be held in place. That silence is precisely where the coercivity law enters, and the proposition we advance is accordingly one of insufficiency rather than of equivalence.

Conjecture (Insufficiency of Requisite Variety). A system may satisfy $V_{acc} \geq V_{req}$ and yet fail to persist. The argument is as follows. Satisfying the variety condition establishes that the operative geometry supports the required number of distinguishable configurations. It establishes nothing about the cost of holding a configuration against perturbation, which is what defect rigidity measures and what the enactment integral accumulates. Where the enactment tension falls short of the threshold that the Saturation-Rigidity Principle requires, perturbation is not resisted at the rate needed to restore the configuration; the effective rank of the Fisher matrix on the operative manifold falls, and V_{acc} falls with it, so that a system which satisfied the variety condition before perturbation may fail it afterwards. Variety, on this reading, is not a stock a regulator possesses but a quantity it must continuously hold, and the coercivity law states the cost of holding it. This is the formal content of the informal observation that a system may have the right shape and still be unable to pay for it.

The conjecture is stated as such, and its two open points should be visible rather than buried. It presumes an identification between the integrated informational potential of the coercivity law and the spectrum of the Fisher matrix on a parametric subfamily; these are related but not identical objects, the former being an integral over the manifold and the latter a local spectral quantity, and the step between them is not taken here. It presumes further that the loss of effective rank under sub-threshold perturbation is monotone in the shortfall, which is plausible from the structure of the admissibility gap but is not proved. What the conjecture does supply is a falsifiable statement in place of an analogy: if a system that satisfies the variety condition is shown to retain its accessible variety under perturbation while its enactment tension lies below threshold, the proposed connection fails for that domain. We would regard that as failing.

The LLM example of Section 1 can now be stated more sharply. An LLM's I_p (the integrated informational potential of its operative field) is, in proxy terms, the richness and precision of the regulatory information it holds in place during inference. Its c (the coherence conductance) is the speed at which informational structure locks into place during a forward pass. Both are non-trivial and grow with scale and training; an LLM is not a trivial system. What the framework predicts, however, is that $I_p c^2$ for an LLM-in-isolation sits below the coercivity threshold, because the system has no operative manifold of its own on which rigidity can accumulate across inference calls: each call begins from a frozen weight state, produces a transient activation pattern, and discards it. There is no K-field internal to the LLM that supplies long-horizon boundary conditions across calls; those boundary conditions are supplied externally by the prompt, the platform, the training pipeline, and the human user.

The two-condition framework makes a specific, falsifiable claim: no amount of parameter scaling alone will cross the coercivity threshold for an LLM, because the threshold is a property of the operative manifold's accumulated rigidity, not of the parameter count. Crossing the threshold would require architectural changes that let rigidity accumulate on a persistent operative manifold (persistent memory, persistent goal-state, persistent self-model) rather than being re-initialized on each call. Whether such architectures are feasible is an open engineering question. The framework's contribution is to name the property they would need to possess, and to give it a formal criterion. An LLM has the variety (often vastly more than requisite), but lacks the slack. The margin between its actual $I_P c^2$ and the coercivity threshold is closed not by the model itself but by its scaffolding.

Two qualifications belong with this claim. The first concerns what is not offered. No illustrative values, even to an order of magnitude, are computed for I_p or c for a language model, and this follows from the constraint set out above rather than from an omission. Producing such a figure would require exactly the conversion between formal quantities and measurable rates that BC3 withholds, and supplying one without that conversion would yield a number with no defensible units, which is the appearance of rigor rather than rigor. The claim advanced here is accordingly comparative and architectural. It concerns the absence of a persistent operative manifold on which rigidity could accumulate, and it stands or falls on that architectural feature rather than on any computed threshold. The second qualification concerns where the system boundary is drawn. The argument is about a language model in isolation. A model coupled to prompt engineering, retrieval, in-context learning, tool use, and a human interlocutor is a different system with a different boundary, and nothing established here shows that the coupled system fails the two conditions. The framework's own account, indeed, attributes the coupled system's coherence to precisely those external boundary conditions. Whether the enlarged system satisfies the conditions is a separate question, and an interesting one, which the present analysis does not address.

The CPC-Telegrapher Theorem (T^E_0, complete derivation conditional on the entity-collapse step; see Appendix B.4) establishes that the three-step CPC construction (rigidity lock, entity-collapse, and monodromy reset) generates the telegrapher's equation [48–52]

$$\partial_{tt}\rho + 2\lambda\,\partial_t\rho = c^2\,\partial_{xx}\rho$$

This is the common ancestor of three limiting cases, each of which projects across quantum, classical, and social/organizational domains by the same freezing logic as Table 2. The symbol λ here denotes the damping rate of the telegrapher's equation, and is distinct from the autopoietic ratio J/I of Section 2. Table 3 summarizes the structure.

Table 3. CPC derivations: the telegrapher's equation and its limiting cases.

Equation	Quantum Domain	Classical Domain	Social/Organizational Domain
Telegrapher's equation	Quantum probability-density propagation before the long-time limit	Signal propagation in transmission lines; diffusion with inertia	Information propagation with memory; ballistic transmission + Poisson disruption
Schrödinger equation (long-time limit, Nelson identification $D = \hbar/2m$)	Non-relativistic quantum time evolution	Classical diffusion ($\partial_t \rho = D\, \partial_{xx} \rho$) as overdamped limit	Social diffusion of norms/practices; organizational learning as I-field diffusion
Dirac equation (chiral extension)	Relativistic spin-½ particles, antiparticles	Chiral diffusion with broken symmetry (active matter)	Institutional momentum; constitutional rigidity as a mass-like term

https://doi.org/10.3390/systems14091132

The Schrödinger equation is arguably more structurally significant than the Einstein equation, in two respects (again conditional on the externally-borrowed Nelson identification): it is dynamic, governing time evolution rather than a rest-frame equivalence, and it is domain-transcendent through the telegrapher's equation, the same CPC path recovers classical diffusion and Dirac elsewhere, and projects further into social and organizational domains via Table 2's freezing logic. The derivation of the telegrapher's equation from the CPC reset (the entity-collapse and monodromy-reset steps) is a complete derivation ($\mathrm{T^E}_0$; see Appendix B.4) conditional on the entity-collapse step. That step's explicit Berezin–Pfaffian construction is itself complete only for block-diagonal/normal operators ($\mathrm{T^E}_0$); in the general non-normal case, it remains conjectural and requires independent verification ($\mathrm{T^E}_2$; Appendix B.4), so this part of the claim is unconditional only within the block-diagonal case. What remains conditional is the Nelson stochastic-mechanics identification, where it is invoked to connect the telegrapher's equation to quantum mechanics. Specifically, it remains an externally borrowed result rather than one derived within FGIR ($\mathrm{T^E}_1$).

9. Derivation Summary

The paper's derivation chain runs in seven steps. (1) The SRP (Section 3) imposes conditions S1–S3 on any persisting structure; the Łojasiewicz–Simon gradient inequality [53–56] discharges S3 as a theorem, and the Stability of Saturation lemma shows that persistence implies saturation. The cone proposition of Section 3 then fixes the Lorentzian signature of the projection. (2) On the weighted information manifold, the admissible protoenergy functional is minimized under the Saturation Assumption, discharged for compact analytic families by the Analytic Saturation proposition. (3) The Fundamental Law $E = \mathcal{F}c^2$ follows by squeezing the protoenergy between an upper and a lower bound on the saturation set. (4) Specializing the weight to Fisher information density via the embedding $I \to I_\mathrm{p}$ yields the FGIR coercivity law $E = I_\mathrm{p}c^2$. (5) The freezing projection π maps the incorporeal law to a corporeal domain; the T-BRIDGE theorem is a conditional equivalence on BC1–BC3, with the residual conjectures (Lorentzian signature, time-translation symmetry, strict hyperbolicity, functoriality) each conditionally sketched. (6) The corporeal projections (Section 8) comprise twelve direct-freezing rows plus three CPC derivations; the CPC-Telegrapher Theorem is a complete derivation ($\mathrm{T^E}_0$) conditional on the entity-collapse step, whose explicit construction is $\mathrm{T^E}_0$ for block-diagonal/normal operators and remains conjectural in the general case ($\mathrm{T^E}_2$; Appendix B.4), and the Nelson stochastic-mechanics identification linking it to quantum mechanics remains externally borrowed ($\mathrm{T^E}_1$). (7) The status summary consolidates: the core theorem is proven ($\mathrm{T^E}_0$); the CPC-Telegrapher Theorem is proven ($\mathrm{T^E}_0$) conditional on the entity-collapse construction, which is $\mathrm{T^E}_0$ in the block-diagonal case and $\mathrm{T^E}_2$ (conjectural) in the general case; the bridge to physics is conditional ($\mathrm{T^E}_1$); the empirical program is conjectural ($\mathrm{T^E}_2$) and awaits data. What remains open, in short, is the bridge to physics and empirical testing; both are addressed directly in Section 10.

10. Conclusions and Research Program

The Constitutional Persistence Cycle (CPC) mechanism's telegrapher-equation derivation is a complete derivation ($\mathrm{T^E}_0$) conditional on the entity-collapse step: the CPC-Telegrapher Theorem establishes that, given that collapse, the three-step construction (rigidity lock, entity collapse, monodromy reset) yields the telegrapher's equation and its Schrödinger/diffusion/Dirac limits. The explicit Berezin–Pfaffian construction of the collapse step is itself complete for block-diagonal/normal operators ($\mathrm{T^E}_0$) but remains conjectural in the general non-normal case ($\mathrm{T^E}_2$; Appendix B.4), and the Nelson stochastic-mechanics identification linking the CPC output to quantum mechanics remains externally borrowed ($T^E{}_1$). The bridge to physics is conditional ($T^E{}_1$): T-BRIDGE establishes that $E = I_p c^2$ and $\hat{E} = \hat{m}\hat{c}^2$ are equivalent representations of a single constitutional persistence law if and only if BC1–BC3 hold, and those are working conditions adopted at R(6), not theorems proven here. The ontological status of the four-field architecture is that of a working commitment grounded in the framework's axiomatic basis (A1–A10): the paper derives the coercivity law from the SRP and weighted-metric geometry conditional on that axiomatic commitment. It leaves the ontology's own derivability (the Minimality Conjecture) as a separate program.

Four conjectures would (if discharged) convert T-BRIDGE from conditional to unconditional. The first, R6-SIG, establishes the four-dimensionality of the projected metric and so discharges the residual half of BC1; its signature is fixed by the cone proposition of Section 3. The second, NOETH-SYM, establishes the time-translation symmetry of the frozen field equation as a Noether theorem and discharges BC2. The third, HYPER, establishes the strict hyperbolicity of the frozen field equation via principal-symbol computation and discharges BC3. The fourth, a functoriality-discharge conjecture, establishes the monoidal functoriality of π and reduces T-BRIDGE to BC3 alone. A fifth, the Minimality Conjecture for the four-field ontology, would convert the ontology from a modelling choice into a derived consequence if resolved.

FGIR treats information, not matter, energy, or spacetime, as the primary generative substrate of reality: what we ordinarily call physical objects are stabilized patterns of information that have undergone enaction—the passage from a dispositional, potential state in the J-field to a structured, enacted state in the I-field, and eventually to a corporeal manifestation in the C-field. $E = I_p c^2$ is the formal quantification of that process. This connects to the embodied-mind tradition of [2] and its extension by Di Paolo [57], who showed how adaptive agency emerges from structural coupling as an organism actively maintains its viability boundaries; FGIR extends this from the biological to the informational domain generally, treating the SRP as the formal precondition of persistence in any domain and the coercivity law as its quantitative consequence. FGIR does not claim to reduce physics to information theory: what T-BRIDGE establishes, conditionally, is that the informational and physical descriptions of constitutional persistence are inter-derivable (the same equation in two ontological registers) given the freezing correspondence.

This paper's primary contribution is the citable derivation of the coercivity law within the FGIR program: the first self-contained, precisely conditional derivation showing that $E = I_p c^2$ follows from the Saturation-Rigidity Principle, the Łojasiewicz–Simon gradient inequality, the Bakry–Émery curvature-dimension condition [25], and Chentsov's uniqueness theorem [26], given the Saturation Assumption and the Yolles ontological embedding as working postulates. The paper completes, as a direct sequel, the architectural derivation of [1]: where that paper showed when recursive closure at R(3) is structurally possible and why it produces a 2–3–2–3 trait parity, this paper shows what that closure costs and what threshold must be crossed for a merely adaptive CAS to become a fully autopoietic CAAS. Within the FGIR architectural map, this paper occupies the derivation-relay layer between the framework's axiomatic architecture and its higher applied layers (the Constitutional Persistence Cycle, the Wuxing reconstruction, and the Beer VSM application [58]), all of which invoke $E = Ipc^2$ as a structural law that this paper now grounds rather than stipulates.

The coercivity law makes testable predictions that are the subject of ongoing work, not a confirmed empirical law. The companion Enactment Tension paper [28] provides a minimal falsifiable footprint for critical-slowing-down predictions, in the sense established for critical transitions generally [59], applicable to non-compact long-memory domains (biology, cognition, social systems, computation) without requiring prior calibration or bridge-condition resolution. The full proxy-tracking and resonance-signature protocol is developed in [13]. A pre-committed Bloch-sphere/transmon-qubit anharmonicity test provides a Popperian test of the Fisher–Rao spectral structure on the simplest non-trivial compact manifold; a spectral-gap test on SU(3) is a candidate for future work. The dimensional-identification problem for I_p in non-physical domains now has a formal resolution criterion (the Proxy Tracking Theorem of [13]) but that criterion has been carried through in full for only one domain to date: the geopolitical domain, via the Iran–Israel/US study [16], which is identified [13] as the framework's first structured application of the four-condition protocol. The organizational domain has a qualitative precedent (Rautakivi and Yolles' ASEAN diagnosis [44,45]) but not yet a quantitative one. The remaining rows of Table 2 have neither. The cross-domain empirical program cannot proceed quantitatively at scale until this protocol has been carried out, domain by domain, beyond the single case where it currently stands—this is the chief barrier to empirical testing outside physics, and is stated here rather than glossed over. Readers seeking to apply the framework to a new domain should begin with the protocol of [13] and the falsifiable footprint of [28], both available via their Zenodo DOIs.

The methodological gain of this bounded program is structural transparency: each residual conjecture names a precise mathematical target addressable by standard techniques—recursion-theoretic and differential-geometric analysis, variational analysis, PDE analysis of the frozen Euler–Lagrange equation, and principal-symbol computation. None is trivial, but each is well-posed, and their cumulative discharge would convert T-BRIDGE from a conditional equivalence into an unconditional theorem. The coercivity law $E = I_p c^2$ is, as it stands, a theorem; its connection to $\hat{E} = \hat{m}\,\hat{c}^2$ is a precisely conditional equivalence whose residual conjectures are sharply bounded and structurally necessary. Read alongside [1], the two papers together give the FGIR program's account of third-order cybernetic closure: where it occurs, why it takes the trait-structure it does, and what it costs.

Three results are established as theorems ($T^E{}_0$) within the FGIR framework, conditional on the framework's axiomatic basis (A1–A10). First, the Fundamental Law $E = \mathcal{F}c^2$ on a general information manifold, derived from the Saturation-Rigidity Principle and the Bakry–Émery weighted-metric geometry. Second, the FGIR-specific coercivity law $E = I_p c^2$, derived via the Yolles ontological embedding and Chentsov's coarse-graining uniqueness theorem. Third, the CPC-Telegrapher Theorem, which derives the telegrapher's equation as the common ancestor of the Schrödinger, diffusion, and Dirac equations. Together with [1], these results give a two-condition criterion for autopoietic closure: architectural sufficiency (R(3) fractal-seed closure) and corporeal viability (the coercivity threshold). The criterion is quantitative, testable, and domain-general.

Four things remain open and are stated here so that no claim is presented as stronger than it is. First, the four-field ontology itself is an axiomatic commitment of the FGIR framework, not a derived theorem; the Minimality Conjecture (Table 4, $T^E{}_1$) would convert it from commitment to consequence if discharged. Second, the bridge to physics (the equivalence of $E = I_p c^2$ and to $\hat{E} = \hat{m}\,\hat{c}^2$) is conditional ($T^E{}_1$) on the three bridge conditions BC1–BC3, which are adopted at R(6) closure rather than proven here; four residual conjectures (R6-SIG, NOETH-SYM, HYPER, functoriality-discharge) would discharge them if resolved. Third, the Nelson stochastic-mechanics identification that links the CPC output to quantum mechanics is an externally borrowed result, not derived within FGIR. Fourth, the empirical program is conjectural ($T^E{}_2$): the coercivity law makes testable predictions, but only one domain (the geopolitical, via [16]) has been carried through the four-condition resonance protocol in full. Of the remaining rows of Table 2, the biological, cognitive, social, and organizational rows are $T^E{}_2$-derived projections awaiting that empirical test. The organizational row additionally has candidate C-field indicators proposed here, heuristic in the sense set out in Section 8, and a qualitative empirical precedent (Rautakivi and Yolles' ASEAN diagnosis), as detailed in Section 8. Only the condensed matter, the Bose–Einstein condensate, and cosmological rows remain heuristic formal isomorphisms that have not yet been derived. The cross-domain empirical program is the chief barrier to applying the framework quantitatively outside physics, and is stated here rather than glossed over.

Table 4. Consolidated epistemic-tier summary. $\mathrm{T^E}_0$ = proven; $\mathrm{T^E}_1$ = conditional/sketch-proven; $\mathrm{T^E}_2$ = conjectural/awaiting empirical test.

Claim	Tier
Fundamental Law $E = \mathcal{F}c^2$	$\mathrm{T^E}_0$ (compact analytic manifolds)
Analytic Saturation proposition	$\mathrm{T^E}_0$
Łojasiewicz–Simon theorem	$\mathrm{T^E}_0$
Stability of Saturation lemma	$\mathrm{T^E}_0$ (complete proof)
T-BRIDGE (freezing equivalence)	$\mathrm{T^E}_1$ (conditional on BC1–BC3)
R6-SIG conjecture (dimensional reduction)	$\mathrm{T^E}_2$ (sketch proof)
NOETH-SYM conjecture	$\mathrm{T^E}_1$ (sketch proof)
HYPER conjecture	$\mathrm{T^E}_1$ (sketch proof)
CPC-Telegrapher construction	$\mathrm{T^E}_0$ (conditional on entity-collapse step; construction is $\mathrm{T^E}_0$ block-diagonal, $\mathrm{T^E}_2$ general case—Appendix B.4)
Schrödinger structural-significance claim	$\mathrm{T^E}_1$ (conditional on Nelson identification)
Heuristic and $\mathrm{T^E}_2$ corporeal-projection rows (Table 2)	$\mathrm{T^E}_2$/heuristic
Empirical program	$\mathrm{T^E}_2$ (awaiting data)
Minimality conjecture (four-field ontology)	$\mathrm{T^E}_1$
Bakry–Émery bound (Standing Assumption)	$\mathrm{T^E}_0$ (standard theory [25], applied to the FGIR manifold)
Compact sublevel sets/closure under convergence	$\mathrm{T^E}_0$ (compact analytic); $\mathrm{T^E}_1$ (non-compact)
Rest-mass identification (BC2)	$\mathrm{T^E}_1$ (modelling postulate)

Author Contributions: Conceptualization of the coercivity law C.-K.L., and its Varela embedding M.Y.; formal analysis, M.Y. and C.-K.L.; writing of original draft, M.Y.; writing of review and editing, M.Y. and C.-K.L. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The supporting mathematical material referred to in the text, including the full proofs, the bridge architecture and the test protocols, is available in the companion working paper deposited at https://doi.org/10.5281/zenodo.21592312 (accessed on 2 September 2026).

Acknowledgments: The authors acknowledge the use of a cybernetics-archivist ensemble of large language model systems, employed in an archival and editorial capacity for cross-domain synthesis, terminological consistency checking, and iterative structural critique, governed throughout by the theoretical commitments of the FGIR framework and subject to authorial oversight at every stage. Responsibility for all theoretical claims, mathematical derivations, and interpretive conclusions remains exclusively with the authors. A longer version of this work is available as a working paper [40], which carries the full proofs of the lemmas, propositions and theorems stated here, the

bridge architecture and its dependency graph, the triple-adjunction proof strategy for the Minimality Conjecture, and the pre-committed test protocols, together with additional cross-domain applications. An associated enactment study extends the framework's application to further domains [28].

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A. Notation

Table A1. Core notation used in this paper.

Symbol	Meaning
M, (M, g, $e^{-I}d\mu$)	Information manifold; Bakry–Émery weighted Riemannian manifold
$I(x)$	Information potential (abstract scalar field)
I_p	Integrated informational potential (FGIR-specific embedding of I)
c	Global coercivity constant (coherence conductance)
E	Admissible protoenergy/enactment tension
$\mathcal{F}$	Protomass
A*	Saturation set
J, I, K, C	Dispositional, operative, sustentative, and corporeal fields
π	Freezing projection
$\hat{m}$, $\hat{c}$, $\hat{E}$	Corporeal (frozen) rest mass, signal speed, and energy
BC1–BC3	Bridge conditions (dimensional reduction; constitutional identification; energy preservation)
R(n)	n-th order cybernetic recursion/closure level
SRP; S1–S3	Saturation-Rigidity Principle and its three conditions
CAS, CAAS	Complex Adaptive System; Complex Adaptive Autopoietic System
CPC	Constitutional Persistence Cycle
$T^E{}_0$, $T^E{}_1$, $T^E{}_2$	Epistemic tiers: proven; conditional; conjectural (see Table 4)

Appendix B. Mathematical Detail

This appendix collects the differential-geometric and information-geometric details that support the derivation of Section 4, so that the main text can carry only the logical skeleton. Readers who require the full proofs should consult this appendix together with the working paper [40], which carries them in full.

Appendix B.1. The Bakry–Émery Framework

The base geometry is a weighted Riemannian space $\left(M, g, e^{-I}d\mu\right)$, where g is an ordinary Riemannian metric and $e^{-I}d\mu$ is the Bakry–Émery volume form weighted by the information potential I [25]. Three structures follow immediately. The weighted Laplacian $\Delta_I = \Delta - \nabla I \cdot \nabla$ has carré du champ $\Gamma_I(f) = \parallel \nabla f \parallel^2$—an algebraic identity that fixes the squared-gradient exponent $\alpha = 2$ in the admissible protoenergy functional as a structural consequence of the calculus, not a modelling choice. The iterated carré du champ satisfies the curvature–dimension condition $\mathrm{Ric} + \nabla^2 I \geq Kg$ for some $K > 0$ (the $\mathrm{CD}(K, \infty)$ condition), supplying the spectral-gap input for the Stability of Saturation lemma. And the weighted Poincaré and log-Sobolev inequalities have right-hand sides equal to the Dirichlet form, or equivalently, the Fisher information, when f is a score function.

Appendix B.2. Chentsov's Theorem and Fisher–Rao Structure

Where a statistical subfamily $\{\rho(\cdot \mid \theta) : \theta \in \Theta\}$ is parametrized, the Fisher–Rao metric

$$g_{ij}(\theta) = \int [\partial_i \log\rho(\cdot \mid \theta)][\partial_j \log\rho(\cdot \mid \theta)]\, \rho(\cdot \mid \theta)\, dx$$

is, by Chentsov's theorem [26,38,39], the unique coarse-graining metric (up to a positive scalar) invariant under all Markov morphisms. Its connection to SRP-S2 is mediated by the Kullback–Leibler divergence, whose local quadratic expansion

$$D_{\mathrm{KL}}(\rho(\cdot \mid \theta) \parallel \rho(\cdot \mid \theta_0)) = \frac{1}{2}\, g_{ij}(\theta_0)\, \Delta\theta_i\, \Delta\theta_j + O(\mid \Delta\theta \mid^3)$$

is exactly the squared-defect structure S2 demands. The Fisher–Rao metric is not the base metric of the incorporeal manifold; it enters only as the coarse-graining variational metric on whichever parametrized subfamilies are locally in view (for example, the qubit Bloch-sphere test of [40]).

Appendix B.3. The Derivation in Detail

Lemma A1. (Coercive upper bound)**.** *If c is finite, then*

$$\int I \parallel \nabla_g I \parallel^2 \leq c^2\, \mathcal{F}.$$

Proof. The pointwise gradient of I cannot exceed c anywhere on the manifold, so the integrand is bounded above by Ic^2 and the integral by $c^2 \int I\, d\mu_g = c^2 \mathcal{F}$. □

Lemma A2. (Lower bound via saturation sets)**.** *For measurable A,*

$$\int_A I \parallel \nabla_g I \parallel^2 \geq \operatorname*{ess\,inf}_A I \cdot \int_A \parallel \nabla_g I \parallel^2 .$$

Proof. On A, I is bounded below by its essential infimum, which factors out of the integral. □

Saturation Assumption. There exists an energy-minimizing metric g^* and a measurable saturation set A^* of positive measure satisfying (i) gradient saturation, (ii) potential uniformity, (iii) energy concentration, and (iv) mass concentration. The four conditions and their logical independence are stated in [40].

Analytic Saturation proposition. For compact analytic statistical families, the admissible protoenergy functional is real analytic, and its critical set contains a nonempty saturation stratum A^* satisfying (i)–(iv). This discharges both the Saturation Assumption and the real analyticity hypothesis underlying the Łojasiewicz–Simon theorem. This theorem guarantees that for any real-analytic energy functional, a gradient flow that gets close to a critical point must converge to it, with the theorem providing the precise rate at which the convergence occurs.

Fundamental Law. Under the above,

$$E = Fc^2.$$

Proof. Upper bound: Lemma A1 gives $E \leq c^2 F$. Lower bound: by energy concentration,

$$E = \int_{A^*} I \; \| \nabla_{g^*} I \|^2 \; d\mu_{g^*}.$$

Conditions (i)–(ii) give $I \; \| \nabla I \|^2 = \bar{I} \, c^2$ on A^*, so

$$E = \bar{I} \, c^2 \, \mu_g(A^*) = c^2 F.$$

□

Appendix B.4. The CPC Construction

The Constitutional Persistence Cycle (CPC) generates the telegrapher's equation through three steps: a Łojasiewicz rigidity lock (converting gradient-flow convergence into a discrete locked state); a Berezin–Pfaffian entity-collapse step (converting the locked state into a persistent random walk); and a Picard–Lefschetz monodromy reset (closing the cycle). Given that collapse, the CPC-Telegrapher Theorem (that these three steps generate the telegrapher's equation) is a complete derivation ($\mathrm{T^E}_0$). The explicit algebraic construction of the entity-collapse step is itself graded rather than uniformly $\mathrm{T^E}_0$: the Berezin-integral/Pfaffian calculus and the collapse definition are complete for block-diagonal/normal operators ($\mathrm{T^E}_0$); the identification of the collapse with the spectral phase of the rigidity operator in the general non-normal case remains conjectural and requires independent verification of the underlying polar decomposition ($\mathrm{T^E}_2$); and the connection between this construction and the persistent-random-walk reading is offered as an interpretation rather than a theorem. The Nelson stochastic-mechanics identification that links the CPC output to quantum mechanics remains externally borrowed ($\mathrm{T^E}_1$), as flagged in Table 4.

Appendix B.5. $R(6)$ Closure and the Bridge Conditions

$\mathrm{R}(n)$ closure denotes the n-th order recursion loop's satisfaction of the closure conditions appropriate to that level. $\mathrm{R}(6)$ closure (the level at which the 2–3–2–3 parity alternation of [1] has completed three full cycles) is where π is defined, because the incorporeal triad has stabilized sufficiently to permit a consistent corporeal projection. $\mathrm{R}(6)$ closure requires: (i) the autopoietic $J{\ulcorner,\urcorner} \leftrightarrow {\ulcorner,\urcorner}I$ loop closed (A^* non-empty, gradient saturation holding); (ii) the autogenetic $K{\ulcorner,\urcorner} \leftrightarrow {\ulcorner,\urcorner}\left(J{\ulcorner,\urcorner} - {\ulcorner,\urcorner}I\right)$ loop closed ($\lambda = J/I$ stable); (iii) the autopraxic $\left(J{\ulcorner,\urcorner} - {\ulcorner,\urcorner}I{\ulcorner,\urcorner} - {\ulcorner,\urcorner}K\right){\ulcorner,\urcorner} \leftrightarrow {\ulcorner,\urcorner}C$ loop closed (π defined); (iv) three full parity cycles completed; and (v) the $R(6)$ fixed-point metric Lorentzian of signature $(3, 1)$, a separate, non-automatic condition, the content of the R6-SIG conjecture.

References

1. Yolles, M. Metacybernetics: Aspect Traits and Fractal Patterns in Higher-Order Cybernetics. *Systems* **2026**, *14*, 496. [CrossRef]
2. Varela, F.J.; Thompson, E.; Rosch, E. *The Embodied Mind: Cognitive Science and Human Experience*; MIT Press: Cambridge, MA, USA, 1991.
3. Schwarz, E.; Aragno, M.; Beck, H.; Matthey, W.; Remane, J.; Chiffelle, F.; Gern, J.-P.; Dubied, P.-L.; Bühle, P. *La Révolution des Systèmes: Une Introduction à l'approche Systémique*; Secrétariat de l'Université, DelVal: Neuchâtel/Cousset, Switzerland, 1988.
4. Schwarz, E. A generic model describing the complexification and autonomization of natural systems, and its epistemological consequences. In *Advances in Systems Studies*; Lasker, G.E., Ed.; IIAS: Windsor, ON, Canada, 1993; pp. 37–43.
5. Schwarz, E. A metamodel to interpret the emergence, evolution and functioning of viable natural systems. In *Cybernetics and Systems '94*; Trappl, R., Ed.; World Scientific: Singapore, 1994; pp. 1579–1586.
6. Schwarz, E. Autogenesis. Available online: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3826203 (accessed on 28 August 2026).
7. Holland, J.H. Complex adaptive systems. *Daedalus* **1992**, *121*, 17–30.
8. Holland, J.H. *Hidden Order: How Adaptation Builds Complexity*; Addison-Wesley: Reading, MA, USA, 1995.

9. Varela, F.J. *Principles of Biological Autopoiesis*; North Holland: Amsterdam, The Netherlands, 1979.
10. Maturana, H.R.; Varela, F.J. *Autopoiesis and Cognition: The Realization of the Living*; D. Reidel: Dordrecht, The Netherlands, 1980.
11. Yolles, M.; Frieden, B.R. Viruses as living systems—A metacybernetic view. *Systems* **2022**, *10*, 70. [CrossRef]
12. Yolles, M.I. Informational Realism: The Fisher Information Field Theory. *Zenodo* **2025**. [CrossRef]
13. Yolles, M. *The Proxy Narrowed: Fisher-Metric Tracking and Ontological Projection in Fisher-Generative Informational Realism*; Zenodo: Meyrin, Switzerland, 2026. [CrossRef]
14. Yolles, M.; Chiolerio, A. Quantum-informed cybernetics for collective intelligence in IoT systems. *Appl. Sci.* **2026**, *16*, 10. [CrossRef]
15. Yolles, M. The Butterfly Effect in Civilisational Cascades: A Fisher Information Diagnostic of the 2025 "No Kings" Protests. *Zenodo* **2025**. [CrossRef]
16. Yolles, M.I. Psychohistory and the Mathematics of Collapse: Fisher-Generative Informational Realism Applied to the Iran–Israel/US Conflict System. *Zenodo* **2026**. [CrossRef]
17. Yolles, M.; Fink, G. A general theory of generic modelling and paradigm shifts: Part 1–the fundamentals. *Kybernetes* **2015**, *44*, 283–298. [CrossRef]
18. Nelson, R.R.; Winter, S.G. *An Evolutionary Theory of Economic Change*; Harvard University Press: Cambridge, MA, USA, 1982.
19. Dosi, G. Technological paradigms and technological trajectories. *Res. Policy* **1982**, *11*, 147–162. [CrossRef]
20. Arthur, W.B. Competing technologies, increasing returns, and lock-in by historical events. *Econ. J.* **1989**, *99*, 116–131. [CrossRef]
21. Metcalfe, J.S. *Evolutionary Economics and Creative Destruction*; Routledge: London, UK, 1998.
22. Yolles, M. Diagnosing Market Capitalism: A Metacybernetic View. *Systems* **2024**, *12*, 361. [CrossRef]
23. Yolles, M. *Autogenetic Particle Derivation (APD): An Illustration of Recursive Functionality Within Corporeality*; Zenodo: Meyrin, Switzerland, 2026. [CrossRef]
24. Łojasiewicz, S. *Ensembles Semi-Analytiques*; Notes; Institut des Hautes Études Scientifiques: Bures-sur-Yvette, France, 1965.
25. Bakry, D.; Émery, M. Diffusions hypercontractives. In *Séminaire de Probabilités XIX 1983/84*; Azéma, J., Yor, M., Eds.; Lecture Notes in Mathematics; Springer: Berlin, Germany, 1985; Volume 1123, pp. 177–206. [CrossRef]
26. Chentsov, N.N. *Statistical Decision Rules and Optimal Inference*; Nauka: Moscow, Russia, 1972.
27. Yolles, M. *Theory of Autogenetic Particle Derivation (APD): A Structural Meta-Theory for Photon Bifurcation, Mass Genesis, and Consciousness via Fisher Information Geometry*; Zenodo: Meyrin, Switzerland, 2026. [CrossRef]
28. Yolles, M.; Lin, C.-K. *What Does It Cost to Bring Forth a World? Enactment Tension as a Cybernetic Measure of Structural Persistence*; Zenodo: Meyrin, Switzerland, 2026. [CrossRef]
29. Itô, K. On stochastic differential equations. *Mem. Am. Math. Soc.* **1951**, *4*, 1–51. [CrossRef]
30. Øksendal, B. *Stochastic Differential Equations: An Introduction with Applications*; Springer: Berlin, Germany, 2003.
31. Amari, S.; Nagaoka, H. *Methods of Information Geometry*; American Mathematical Society: Providence, RI, USA, 2000.
32. Frieden, B.R.; Soffer, A. Lagrangians of physics and the principle of extreme physical information. *Phys. Rev. E* **1995**, *52*, 2274–2288. [CrossRef] [PubMed]
33. Yolles, M.I.; Frieden, B.R. Autopoiesis and Its Efficacy—A Metacybernetic View. *Systems* **2021**, *9*, 75. [CrossRef]
34. Frieden, B.R. *Science from Fisher Information: A Unification*, 2nd ed.; Cambridge University Press: Cambridge, UK, 2004.
35. Maturana, H.R. Autopoiesis, structural coupling and cognition: A history of these and other notions in the biology of cognition. *Cybern. Hum. Knowing* **2002**, *9*, 5–34.
36. Hornung, B. Sociocybernetics. In *International Encyclopedia of the Social and Behavioral Sciences*, 2nd ed.; Wright, J.D., Ed.; Elsevier: Amsterdam, The Netherlands, 2015; pp. 912–917.
37. Razeto-Barry, P. Autopoiesis 40 years later. A review and a reformulation. *Orig. Life Evol. Biosph.* **2012**, *42*, 543–567. [CrossRef] [PubMed]
38. Campbell, L.L. An extended Čencov characterization of the information metric. *Proc. Am. Math. Soc.* **1986**, *98*, 135–141. [CrossRef]
39. Ay, N.; Jost, J.; Lê, H.V.; Schwachhöfer, L. *Information Geometry*; Springer: Cham, Switzerland, 2017. [CrossRef]
40. Yolles, M.I.; Lin, C.-K. *The Coercivity Law of Enaction Within Fisher-Generative Informational Realism*; Zenodo: Meyrin, Switzerland, 2026. [CrossRef]
41. Landauer, R. Irreversibility and heat generation in the computing process. *IBM J. Res. Dev.* **1961**, *5*, 183–191. [CrossRef]
42. Yolles, M.; Lin, C.-K. *Enactment Tension Across Disciplinary Frames: A Generative Informational Realism Framework*; Zenodo: Meyrin, Switzerland, 2026. [CrossRef]
43. Hyllus, P.; Laskowski, W.; Krischek, R.; Schwemmer, C.; Wieczorek, W.; Weinfurter, H.; Pezzè, L.; Smerzi, A. Fisher information and multiparticle entanglement. *Phys. Rev. A* **2012**, *85*, 022321. [CrossRef]
44. Yolles, M.; Rautakivi, T. Diagnosing Complex Organisations with Diverse Cultures—Part 1: Agency Theory. *Systems* **2024**, *12*, 8. [CrossRef]
45. Rautakivi, T.; Yolles, M. Diagnosing Complex Organisations with Diverse Cultures—Part 2: Application to ASEAN. *Systems* **2024**, *12*, 107. [CrossRef]

46. Ashby, W.R. *An Introduction to Cybernetics*; Chapman & Hall: London, UK, 1956.
47. Conant, R.C.; Ashby, W.R. Every good regulator of a system must be a model of that system. *Int. J. Syst. Sci.* **1970**, *1*, 89–97. [CrossRef]
48. Goldstein, S. On diffusion by discontinuous movements, and on the telegraph equation. *Q. J. Mech. Appl. Math.* **1951**, *4*, 129–156. [CrossRef]
49. Kac, M. Foundations of kinetic theory. In *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability*; University of California Press: Berkeley, CA, USA, 1956; Volume 3, pp. 171–197.
50. Kac, M. A stochastic model related to the telegrapher's equation. *Rocky Mt. J. Math.* **1974**, *4*, 497–510. [CrossRef]
51. Nelson, E. Derivation of the Schrödinger equation from Newtonian mechanics. *Phys. Rev.* **1966**, *150*, 1079–1085. [CrossRef]
52. Nelson, E. *Quantum Fluctuations*; Princeton University Press: Princeton, NJ, USA, 1985.
53. Simon, L. Asymptotics for a class of non-linear evolution equations, with applications to geometric problems. *Ann. Math.* **1983**, *118*, 525–571. [CrossRef]
54. Simon, L. *Theorems on Regularity and Singularity of Energy Minimizing Maps*; Birkhäuser: Basel, Switzerland, 1996.
55. Jendoubi, M.A. A Simple Unified Approach to Some Convergence Theorems of L. Simon. *J. Funct. Anal.* **1998**, *153*, 187–202. [CrossRef]
56. Chill, R. On the Łojasiewicz–Simon gradient inequality. *J. Funct. Anal.* **2003**, *201*, 572–601. [CrossRef]
57. Di Paolo, E.A. Autopoiesis, adaptivity, teleology, agency. *Phenomenol. Cogn. Sci.* **2005**, *4*, 429–452. [CrossRef]
58. Beer, S. *The Heart of Enterprise*; Wiley: Chichester, UK, 1979.
59. Scheffer, M.; Bascompte, J.; Brock, W.A.; Brovkin, V.; Carpenter, S.R.; Dakos, V.; Held, H.; van Nes, E.H.; Rietkerk, M.; Sugihara, G. Early-warning signals for critical transitions. *Nature* **2009**, *461*, 53–59. [CrossRef] [PubMed]