

The Qi-4 Pillars

Qi-QUA, Qi-EP, Qi-CON, Qi-HAL

The Quantum-Information Backbone of GDOS

Katsushi Arisaka

*Department of Physics and Astronomy, Electrical and Computer Engineering,
University of California, Los Angeles, Los Angeles, CA 90095, USA*

Contact: arisaka@physics.ucla.edu

Abstract

Everything that stays itself has to pay for it. A cell, a flame, a spectrometer, a galaxy: each holds a pattern in place by taking energy in at one boundary and letting it out at another. Physics has been able to price the energy for a century and a half. It has never had a price for the information.

This paper sets one. On the boundary of the six-dimensional bulk there are four currents rather than three — energetic, electric, baryonic, informational — and an elementary event does not spend them independently. It spends all four at once, in one fixed proportion, so that four ledgers are one ledger read in four currencies. The exchange rate between energy and information is a single constant, $Q_B \approx 0.1555$ eV per qubit. It is not fitted: it is the Rydberg energy divided by a ledger integer fixed upstream, scaled by an anomaly coefficient, with nothing adjustable in the chain.

Life is where it can be checked, because life runs on one molecule. The bond ATP breaks carries two of these quanta, and $2Q_B = 30.0$ kJ/mol stands against a measured 30.5 ± 0.5 — agreement to 1.6 percent.

Three pillars follow. One shows that least time in optics, maximum entropy production in thermodynamics and the free-energy principle in neuroscience are one principle read three ways. One predicts a $1/f$ band bounded by *two* knees rather than one unbounded slope, and says what sets each end — a claim that can fail where a bare power law cannot. One shows that anything holding a memory longer than a single act-and-look cycle falls into a definite architecture, whatever it is made of. Each pillar is handed an experiment that can end it.

Together the four are the certificates carried by the configuration where the geometric descent halts. They describe that configuration; they are not the instrument that found it; the converse is a theorem where the landscape has one floor, and open elsewhere.

Executive Summary

Four questions about anything that lasts. Ask what a bacterium, a flame, a market and a brain have in common and the honest answer is short: each is open to what surrounds it, and each keeps being itself anyway. The companion papers of this program supply the stage such a thing stands on — a constitutional charter of four axioms and seven postulates (Arisaka, 2026, A1-Constitution), a six-dimensional geometry whose two extra directions are a scale and a phase rather than further places to go (Arisaka, 2026, A2-GG6), proved free of internal inconsistency (Arisaka, 2026, A3-BI6), and the law describing how any of it looks to an observer who is inside rather than outside (Arisaka, 2026, A5-GDOS). The companion to this one supplies the engine that drives a system downhill and stops it (Arisaka, 2026, A6-GRIP). What none of them supplies is the accounting for what such a system *knows*. This paper asks four questions about anything that persists — what does one elementary event cost, which way does the system go, what should it look like across scales, and how does it remember — and answers each as a theorem with an experiment attached.

The first question: what does one event cost? The answer is that it costs all four currencies at once. Energy, electric charge, baryon number and information leave the bulk through a single channel, so an elementary boundary event does not range over a four-dimensional space of possible bills; it lies on a one-dimensional ray inside it (Qi-QUA, Theorem T-QUA.1). Written for a laboratory, that becomes the Arisaka equation $J_E/Q_B = J_e = J_B = J_{Qi}$: count the events in any one column and you have counted them in all four. Two cautions travel with it. The equality is between normalized event counts, not an assertion that one electron and one bit are the same object; and it forbids something definite — there is no elementary event that carries information while carrying no energy, no charge and no matter, which is the precise sense in which this framework has no room for disembodied information.

The constant, and where it can be caught. The exchange rate in that equation is $Q_B = (2/5) E_{Ry}/K$ with $K = 35$: the Rydberg energy, a ledger integer fixed by the upstream Diophantine closure, and the anomaly coefficient of the modified Bianchi identity. Nothing in it was fitted, and the paper records that as a theorem rather than as a hope — Q_B introduces no new continuous free parameter. Numerically it is 0.1555 eV per qubit, or 15.0 kJ/mol, or about $5.8 k_B T$ at body temperature. Life is where it meets a measurement, because life settled on one energy currency: the phosphoanhydride bond of ATP carries two of these quanta, and $2Q_B = 30.0$ kJ/mol against a measured 30.5 ± 0.5 agrees to 1.6%. A second, independent check sits in the same molecule's machinery. The rotary enzyme that makes ATP pays for each bond with a whole number of protons set by the size of its rotor ring — ten in *E. coli*, eight in oxygen-rich mitochondria, thirteen to fifteen in organisms living on a thinner gradient — and at the measured driving voltage each single proton crossing costs one Q_B . The gear ratio is tuned to the environment; the size of one click is not.

It is worth saying what would refute this, because it is not the number. A calorimeter that disagreed with 0.1555 eV would be arguing with the upstream ledger. What Qi-QUA asserts, and

what an experiment can break, is that *every* elementary event of the channel shares the *same* single conversion factor. One catalytic event whose four-way bookkeeping misses by more than twenty percent ends the pillar.

A word about Landauer. Erasing a bit costs at least $k_B T \ln 2$, and at body temperature Q_B is about eight and a half times that. This looks like a contradiction and is not, because the two statements audit different things. Landauer prices erasure, which is a one-way write along a direction with no way back. Qi-QUA prices carrying, on a circle, where what is stored is a winding that no boundary operation can smooth away. One event, three audits: the circle stores it, the interval seals it, and Qi-QUA bills it.

The second question: which way does it go? Three communities already have an answer, and they do not appear to be the same answer. Optics and mechanics minimize a time; nonequilibrium thermodynamics extremizes entropy production; theoretical neuroscience minimizes a free energy. Qi-EP (Theorem T-EP.1) shows they are three readings of one projected functional, differing only in what is held primary — the path, the currents, or the beliefs. The unification is not verbal: taken as a variational problem in each of the three, the same functional returns the geodesic law, the flux–force law, and the natural-gradient inference rule. It also settles an old quarrel, since maximum entropy production and Prigogine’s minimum entropy production turn out to be conjugate readings of one stationary point rather than rival slogans — one selecting the transport law, the other the state. The same object also fixes where those readings stop: they hold while circulation and gauge motion do no work against the ledger, and the classical counterexamples are exactly the cases where one of them does, which is why the unification is answerable rather than merely tidy. A chemotactic cell exercises all three at once: it follows a path, maintains fluxes, and updates its estimate of where the food is.

The third question: what should it look like across scales? Systems of this kind are known to produce $1/f$ noise, and $1/f$ noise explains very little, because an unbounded power law is compatible with almost anything. Qi-CON (Theorem T-CON.2) makes the sharper claim: the band is finite, and both of its ends are set by quantities the framework names. Below the lower knee, at the slowest surviving relaxation rate, the spectrum flattens; above the upper knee, at the fastest retained microscopic rate, it steepens to f^{-2} . Two knees, not one, and the paper says in advance where to look for each. The width between them is not decoration either: it fixes how many levels deep the memory hierarchy of the next pillar can be.

The fourth question: how does it remember? Once a system holds a trace of its past for longer than a single act-and-look cycle, Qi-HAL (Theorem T-HAL.1) puts its runtime into a definite normal form: a loop running memory to prediction to motion to sensing and back; a phase code; a lattice storing relations rather than snapshots; a ladder of nested timescales. The order of the loop is forced rather than chosen — sensing cannot precede action without the system knowing an observation it has not yet caused, and prediction cannot follow sensing without becoming postdiction. Two consequences are worth stating plainly. Because what is stored is a set of phase differences and

not absolute phases, a memory written on one carrier can be read on a completely different one, which is why a chemical, a neural and a digital system can hold the same content; and because the trace is distributed across relations, a damaged substrate can still reconstruct part of what it held. Memory here is never free: a write is an event, and Qi-QUA bills it.

What this paper has that no companion can supply. The engine imported from GG-A6-GRIP runs in two regimes — descent on a landscape that stays put, and the transient that reorganizes the landscape and hands the next descent its starting point. In cosmology the second kind is exogenous and unrepeatable; the universe does not re-run its own beginnings. The discovery this paper contributes is that a memory-bearing system *internalizes* it. Such a system carries its own dial for reorganizing its landscape, and schedules the reorganizations on evidence: it explores while the evidence is thin, commits when the evidence crosses a threshold, and returns to exploring when the world changes and its predictions start failing. Attention, search, decision, commitment and revision are one normal form, and it is the microscopic mechanism by which a cascade keeps itself going instead of settling once and stopping.

Holding together, and the world seen from inside. Two further structures follow. SeOS names the systems that maintain themselves, and makes the criterion a number rather than a judgment: what comes in through the ports, minus what is dissipated, minus what repair costs, must stay positive. QCI then asks a different question — not whether a system survives but how much of the four-pillar runtime it actually runs — and scores it in a way that cannot be gamed, since a system missing one pillar cannot compensate with excellence in another. On top of these sits the internality layer. A finite system never touches the world; it touches its own boundary, so whatever guides its action has to be reconstructed rather than received. The reconstruction is why motion is part of perception and not an accessory to it: a retina that cannot move is depth-poor, a bacterium that cannot swim is gradient-poor, and the paper makes that precise as a rank condition on what the sensory history can recover. The word used for the internally measured, selected and stabilized part of that reconstruction is a technical one, defined operationally and graded rather than binary, and the paper is explicit that it takes no position on subjective experience.


How far it reaches, and in what sense. The same grammar is then run up a five-story hierarchy — particles, cosmos, chemistry, life, mind — where each story is a fixed-landscape descent and each ascent is a landscape reorganization that opens basins the story below did not have. A floor counts as a floor only if it passes three tests: its collective variables nearly close, it lasts far longer than the relaxation beneath it, and it opens at least one genuinely new exchange channel. The framework's boldest sentence is that life and mind are inevitable, and it is stated in exactly one sense: not that every planet produces them, which is disclaimed, but that in a driven universe whose viable configurations occupy a set of positive measure, a search that keeps returning reaches and stabilizes them with probability one. Higher floors never break lower laws; they change the conditions the lower floors run under.

What is proved here, what is inherited, and what is open. The boundary of a claim matters more than the claim. Proved here: the four pillars as dual-layer theorems, the completeness of the system packet, the internalized regime switch, and the characterization of the ground attractor — in one direction. That direction says the configuration where the descent halts carries all four certificates. The converse, that no other basin carries all four, has a scope and the scope is the companion's. On a stratum the companion proves admissible the converse holds, for a reason that does no independent work: the descent there has exactly one critical point with no direction of descent, so every other critical point fails the halting certificate by the geometry alone. On a stratum whose admissibility is open the converse is open, and the openness is inherited from admissibility, not from anything the pillars say. It is recorded as an open problem with that scope, and it should not be quoted as settled beyond it: asserting it by appeal to a companion whose own argument invokes these same pillars would be a circle, and the two directions are therefore kept apart. Inherited and not re-derived: the charter, the geometry, the boundary law, and the cascade — whose closure carries one named residue of its own, so every construction here that stands on it inherits that conditional status. Named as open rather than claimed: whether the hierarchy is forced to have exactly five floors, and the residual bounds on the inevitability argument. Said plainly about most of the cross-domain material: it is organizational rather than predictive, and the paper marks which rows are which so that a reader cannot mistake a placement for a prediction.

A closing thought. Information has spent a century as physics' guest — indispensable in practice, unpriced in principle, invoked wherever a system had to choose, remember or predict, and quietly excused from the ledger every other quantity had to satisfy. The proposal here is that it was never a guest. It is the fourth column of an account that has always had four columns, and the size of one entry in it is fixed by the same integers that fix how many families of matter there are. If that is right, then the world of currents and free energies and memories that the laboratory sciences actually work in is not the untidy residue left over when fundamental physics has taken what it wants. It is the face the same object turns toward anything finite enough to have to look.


Keywords: GG-Theory, GG-6, GDOS, GRIP, GRIP-In, GRIP-Out, GFF, GEP, CEP, GA, GGA (Ground Geometric Attractor), Qi-4 Pillars, Qi-QUA, Qi-EP, Qi-CON, Qi-HAL, SeOS, QCI, OH, IH, EIC, five-story pagoda, Arisaka Ratio, Arisaka Equation, Arisaka Constant, MePMoS, NHT, HAL, BFFT, internality layer.

Four Information Pillars, and What a Moving System Can Come to Know



What this paper is. GG-A5-GDOS gives the boundary law and GG-A6-GRIP the geometric template that forces it; the present paper supplies the informational half. The geometric half says how an open system moves and what it spends; the informational half says what the spending is denominated in and what the moving system can come to know. Each pillar inherits its constitutional content from GDOS paired with the Gi-4 template above, and its bulk-side content from the 6D gauge sector below; the horizontal arrows are implications, not analogies. Neither claim survives on its own — a reader who takes only the geometry has a runtime with an unpriced column, and one who takes only the information has a price without an engine. In full: Figure 4, p. 56.

One Runtime, Five Floors: Particle, Cosmos, Chemistry, Life, Mind



Pagoda factorization theorem: each higher floor is forced by closure-defect + port-novelty admissibility on the previous.
 Conditional inevitability theorem: in a sufficiently rich driven universe, the higher floors are reached almost surely (ensemble sense).

The same runtime, five times. Particle, cosmos, chemistry, life, mind. Each floor is forced by the one beneath it through the pagoda factorization theorem — closure-defect plus port-novelty admissibility — and within each floor the dynamics is GRIP-In, descent to that floor's ground attractor, while the ascent from one floor to the next is a GRIP-Out event that opens basins the lower floor did not possess. The *shinbashira*, the central pillar, is the OGGEp projector carrying the combined Gi-4/Qi-4 cascade, and it runs through every floor unchanged. Two honesties are on the plate rather than behind it: five is the historically natural sequence and is not proved unique, and the inevitability of the upper floors is conditional on a driven universe with positive viable subsets. In full: Figure 16, p. 201.

Five Predicted Signatures, and What a Failure of Each Would Mean

Where this framework is exposed. Five predicted signatures, each with a test and — the column that gives the table its value — what a failure would mean. A framework this broad can absorb a surprising amount of data unless it says in advance which failures are fatal and which are merely inconvenient, and that is what the third column is for. All five are open in the experimental sense: the prediction is on the page and the verdict is the next decade of measurements. In full: Table 54, p. 228.

Predicted signature	Test idea	What failure would mean
Current-locking constraints	Estimate event-level energy/charge/information budgets in repeated micro-events; test whether one locked ledger explains them.	Persistent inability to define a coherent event ledger would challenge Qi-QUA.
Selector behavior	Perturb pathway options and test whether realized trajectories follow a common PLA/PLT/MEPP/FEPstyle selector.	If no selector family fits across domains, Qi-EP loses force.
Two-knee scale signatures	Measure spectra, burst statistics, and fractal indicators under controlled conditions.	Absence of multiscale diagnostic structure would weaken Qi-CON.
Exploration-to-commitment lag	Time-resolve candidate-state exploration, branch selection, and stable report or action.	If selection/stabilization staging is absent, the DGI $\rightarrow$ GSSB $\rightarrow$ GA pipeline is incomplete.
Motion-conditioned model quality	Restrict action while preserving passive input and test whether internal model quality degrades.	Failure would challenge the action-conditioned observability theorem behind IH.

Overview of Contents

Part I. Entrance and bridge	20
Part II. Runtime constitution imported from GG-A6-GRIP	48
Part III. Qi-QUA	58
Part IV. Qi-EP	77
Part V. Qi-CON	124
Part VI. Qi-HAL	141
Part VII. SeOS and QCI	164
Part VIII. The Qi-4 pillars as the GGA characterization	174
Part IX. Internal Measurement and the EIC Runtime	180
Part X. Emergence and scope	200
Part XI. Assessment	217
Part XII. Conclusions	232
Part XIII. Technical Appendices	242
Part XIV. Common Objections and Responses	294
Part XV. References	321

List of Figures

1	The open-system chain shared across the GG-Theory cohort.	29
2	The twenty-two-paper cohort map.	44
3	Roadmap of the fifteen Parts of GG-A7-Qi4.	46
4	Architecture overview of the four Qi-4 pillars.	56
5	The Arisaka equation as the four-current lock of Qi-QUA (Pillar 1).	59
6	The circle that remembers and the interval that forgets.	74
7	Three centuries, three traditions, one functional.	80
8	One functional read in three gauges.	94
9	Why the entropy readings fail where circulation does work.	110
10	The runtime's three moves on the selector's landscape.	118
11	The two-knee spectrum: a band, not a slope.	126
12	The memory loop, and why its order is forced.	143
13	The maintenance margin, the viability tube, and the second question.	165
14	Four certificates, and the direction that is proved.	178
15	The OH+IH=EIC framework of Part IX.	181
16	The five-story pagoda hierarchy of Part X.	201
17	Hierarchy of program-defined acronyms in the GG-Theory program.	247

List of Tables

1	Why living systems are the natural entry point.	23
2	Claims and non-claims of this paper.	27
3	Division of labor across the foundation papers.	30
4	The imports, their sources and their roles.	32
5	The open-system logic across five domains.	36
6	Why the kernel needs four current sectors.	39
7	An operational audit of the runtime pipeline.	50
8	The four ledgers audited by Qi-QUA.	61
9	Three equivalent readings of the Arisaka Ratio.	64
10	The Arisaka constant in four unit systems.	68
11	The four-line derivation of Qi-QUA.	71
12	Ideal and coarse-grained forms of Qi-QUA.	72
13	Landauer and Qi-QUA answer different questions.	73
14	The action tradition: what is proved, where it stops.	83
15	MinEP: the hypotheses and their counterexamples.	86
16	MEPP: one theorem, one heuristic, one failed proof.	89
17	FEP: four strata, four verdicts.	92
18	The least-time gauge, step by step.	99
19	One least-time law, three centuries apart.	100
20	Five claims, one stationary point: the quarrel settled.	104
21	The entropy-production gauge, step by step.	105
22	What the FEP assumes, and what the gauge derives.	108
23	The inferential gauge, step by step.	108
24	The three gauges read against eight observations.	114
25	The three gauges of one functional.	116
26	The grade of every step in the Qi-EP chain.	116
27	Which gauge is natural in which domain.	119
28	The four signatures of a Qi-CON window.	127
29	One exponent, read three ways.	134
30	What the two spectral knees mean.	137
31	How the earlier layers constrain Qi-CON.	138
32	Expected Qi-CON readouts by domain.	139
33	The five motifs of the memory architecture.	146
34	The memory-prediction loop across four domains.	150
35	Phase-based encoding beyond neuroscience.	151
36	The multiscale hierarchy, level by level.	156
37	The five motifs, system by system.	161
38	Open, packet-bearing, and self-owning systems.	168
39	What each of the four pillar scores measures.	170
40	How packet, ledger, runtime and persistence fit.	171

41	Outer, inner, and the coupling of the two.	182
42	Boundary encoding in four domains.	184
43	A graded continuum, not a binary.	194
44	What this section does and does not claim.	196
45	What the framework predicts here, and where.	196
46	The five floors and what each adds.	205
47	The pagoda analogy, translated.	207
48	Operational criteria for each floor.	209
49	How to apply the framework, step by step.	212
50	What each floor of the hierarchy inherits, and what would refute it.	213
51	How this framework differs from six neighbors.	218
52	What this paper imports, and from where.	219
53	What this paper exports, and to whom.	221
54	The falsification program, signature by signature.	228
55	Summary of abbreviations used in the present paper.	243
56	Glossary of key mathematical symbols.	248
57	The principal equations of the framework, in order.	249
58	The six kinds of attractor the definition allows.	253
59	Unit conversions for the Arizaka constant.	255
60	Ideal event budgets, one to eight quanta.	256
61	A working dictionary for the three gauges.	262
62	A working dictionary of Qi-CON quantities.	268
63	How unfolded data groups into the packet.	271
64	Packet, ledger and persistence, domain by domain.	273
65	Outer and inner holograms, system by system.	276
66	Application cards for four domains.	278
67	The falsifiability matrix, block by block.	279
68	The packet slots, read in chemistry.	281
69	Chemistry-scale diagnostics, claim by claim.	284
70	Neural variables, mapped to the runtime.	286
71	Candidate criteria for a reportable state.	288
72	The four channels, read collectively.	291
73	Notation concordance: the preferred form of each object.	293

Contents

Executive Summary	2
List of Figures	9
List of Tables	10
Preface	19
 Part I. Entrance and bridge	 20
1 Introduction	21
1.1 Why living organisms are the natural entry point for Paper IV-B	21
1.2 From the measurement problem to the universal observation problem	23
1.3 The need for a common runtime	25
1.4 Preview of the new objects introduced in Paper IV-B	26
1.5 Claims and non-claims of the present paper	27
1.6 Roadmap of the full manuscript	27
2 This paper's place in the cohort: one link of the open-system chain	28
3 Recap bridge to GG-A1-Constitution, GG-A2-GG6, GG-A3-BI6, GG-A5-GDOS, and GG-A6-GRIP	28
3.1 Why the recap is structurally necessary	28
3.2 Paper I recap: the GG-Constitution axiom system	30
3.3 Paper II recap: GG-6, BI-6, and the six-dimensional bulk	30
3.4 Paper III recap: GDOS as the boundary foundation	31
3.5 Paper IV-A recap: GRIP, GFF, and SIP	31
3.6 Bridge equation to Paper IV-B	32
4 Universal extension of GDOS to any open system in nature	33
4.1 From the boundary law of Paper III to a universal runtime	33
4.2 Why every scientific observation is constitutionally open	35
4.3 Ports, currents, and the minimal bookkeeping common to all systems	36
4.4 Why the new gauge term is unavoidable	38
4.5 Section summary: what has now been universalized	39
5 The cohort inheritance and the roadmap of the fifteen Parts	40
5.1 The cohort inheritance	40
5.2 Roadmap of the fifteen Parts	43
6 Why none of this is a special case	47
 Part II. Runtime constitution imported from GG-A6-GRIP	 48
7 The imported runtime theorem: GRIP	49
8 GFF: the geometric free-energy functional $GFF = GEP + CEP$	51
9 SIP: System Instance Packet as the Canonical Unit of GDOS	53
10 What had to be true first	55
10.1 What this Part returns to the companion it borrows from	56
 Part III. Qi-QUA	 58
10.2 Historical background (with citations)	59
11 The four ledgers of a generic open system	61
12 From BI-6 to a one-dimensional boundary source ray	62
13 The Arizaka Ratio and its equivalent representations	64
14 From GG-6 scales to the Arizaka Constant	66
15 Natural-unit reduction and the Arizaka Equation	68
16 Ideal equality, coarse-grained inequality, and overhead	71

17	Relation to Landauer and to classical nonequilibrium thermodynamics	72
18	Information and the two fibers: memory on the circle, erasure on the interval . . .	74
19	What it means to put a price on a bit	75
Part IV. Qi-EP		77
20	Overview of PLT, MEPP, and FEP	78
21	The principle of least action and least time: history, applications, and limitations .	79
21.1	From Fermat to Hamilton	79
21.2	The twentieth century: why “least,” and what noise does	81
21.3	Applications	81
21.4	Limitations	82
22	Minimum entropy production (Prigogine): history, applications, and limitations . .	83
22.1	A century of minimum-dissipation theorems	83
22.2	The theorem, exactly	84
22.3	The failed generalization, and what it created	84
22.4	Critiques, counterexamples, and the exact scope	85
23	Maximum entropy production: history, applications, and limitations	86
23.1	From least dissipation to a maximum principle	86
23.2	Paltridge’s climates: the empirical core	87
23.3	The claimed derivation, and where it broke	88
23.4	Scope, and the distinction that prevents most confusion	88
24	The free energy principle: history, applications, and limitations	89
24.1	Three streams, then a fourth	89
24.2	The mathematics that is solid	90
24.3	The contested stratum, and the critiques that located it	91
24.4	The honest summary, and what a derivation owes it	92
25	Constitutional statement of the Qi-EP principle	93
25.1	The contract: five criteria a genuine derivation must meet	94
26	From the 6D bulk action to the projected boundary action	95
26.1	The imports, with their grades	95
26.2	The mathematical tools, briefly	95
26.3	The projection, step by step	96
27	Trajectory gauge: derivation of PLT	98
27.1	The regime, and a worked parallel	99
28	Current gauge: derivation of MEPP	100
28.1	The third reading, and Prigogine’s theorem recovered	103
29	Inferential gauge: derivation of FEP	105
29.1	What the inferential gauge does not assume	107
30	Where the three gauges must fail	108
31	The three gauges in the laboratory	110
31.1	Path selection in a structured medium	111
31.2	Thermodiffusion: the textbook force-side state	111
31.3	Planetary heat transport: the flux side at scale	111
31.4	Jaynes’s two-temperature network: the counterexample as a theorem . . .	112
31.5	Bénard convection: the boundary of both readings	112
31.6	Chemotaxis: the inferential gauge with a checkable rank condition	112
31.7	Neural prediction: the gauge the principle was built for	113
31.8	ATP synthesis: all three gauges in one molecule	113
31.9	What the eight cases share	114
32	The unified Qi-EP theorem: one projected functional, three gauges	115
32.1	The contract, discharged	116
33	DGI, GSSB, and GA inside the Qi-EP landscape	117

34	Why Qi-EP applies to all open systems	118
34.1	Nothing biological was used	119
34.2	The cascade demonstration	120
35	Three sciences, one principle	120
36	Concluding remarks: one landscape, three shadows	122
Part V. Qi-CON		124
36.1	Historical background (with citations)	125
37	Constitutional statement of the Qi-CON principle	125
38	From the 6D bulk to homogeneous boundary kernels	128
39	The leading Qi-CON generator and operational conformality	130
40	Second signature: Lévy statistics from the fractional generator	132
41	Third signature: fractality of trajectories and selected supports	133
42	Fourth signature: $1/f$ spectra with two knees	134
43	Why the Qi-CON exponents are runtime invariants	137
44	Why Qi-CON applies to all open systems	138
45	Why two knees are worth more than one slope	139
Part VI. Qi-HAL		141
45.1	Historical background (with citations)	142
46	Constitutional statement of the Qi-HAL principle	142
47	Why persistent memory plus $U(1)_{Qi}$ is the true entrance condition	147
48	MePMoS: derivation of the causal loop	148
49	NHT: phase-holographic coding of latent geometry	150
50	HAL: the holographic attractor lattice as a memory substrate	152
51	BFFT: the dyadic multiscale protocol	154
52	The GRIP-loop as the exploration-to-commitment switch	156
53	A unified effective action for Qi-HAL	157
54	How Qi-HAL applies to any open system with memory	160
55	Memory as something the world does	162
Part VII. SeOS and QCI		164
56	Why persistence and completeness belong after the full Qi stack	166
56.1	Historical background (with citations)	166
57	SeOS: the class of persistent self-maintaining open systems	166
58	The persistence theorem	167
59	QCI: the Qi-4 Completeness Index	169
60	Integrated view: from runtime constitution to persistence and completeness	170
61	Staying alive, and how much of it you are doing	172
Part VIII. The Qi-4 pillars as the GGA characterization		174
62	The characterization theorem, pillar by pillar	174
63	Four certificates, and what they are not	179
Part IX. Internal Measurement and the EIC Runtime		180
64	OH+ IH= EIC: Outer Hologram, Inner Hologram, and Internal Measurement	180
64.1	Why the runtime still needs a distinction between external projection and internal reconstruction	181

64.2	The Outer Hologram: lawful projection from the 6D bulk to measurable boundary records	183
64.3	The Inner Hologram: internal reconstruction of allocentric structure from partial sensory streams	184
64.4	Why motion and memory are mathematically necessary for the inner hologram	186
64.5	Why the inner hologram must be allocentric rather than a stored sensory movie	187
64.6	The equality $OH + IH = EIC$: precise meaning of the plus sign	188
64.7	Consciousness as internal measurement of the inner hologram	190
64.8	Universality across open systems with memory once $U(1)_{Q_i}$ is present . . .	192
64.9	Claims, non-claims, and empirical consequences	195
64.10	Section summary	197
65	The world one can never reach	198
Part X. Emergence and scope		200
66	Hierarchy and emergence: the Japanese five-story pagoda	201
66.1	Why a universal theory of open systems must become hierarchical	202
66.2	The pagoda factorization theorem	202
66.3	DGI as the pagoda's central pillar (<i>shinbashira</i>)	205
66.4	The five floors in detail	207
66.5	Probabilistic recurrence of high-floor packets in a driven open universe (conditional)	209
66.6	Upward inheritance and downward modulation	210
66.7	Section summary	210
67	Cross-domain applications: chemistry, life, brains, and society	211
67.1	A canonical reduction recipe for any open system with memory	211
68	What each floor is answerable for	212
69	The exportable result: inference derived, not assumed	214
70	Why GDOS applies to all open systems in nature	215
71	One grammar, many sciences	215
Part XI. Assessment		217
72	Discussion: significance, uniqueness, and comparison with existing frameworks . . .	217
72.1	What is new in the present paper	217
72.2	Relation to existing scientific programs	218
73	Cross-Paper Coherence Map	219
73.1	Inward map: imports into GG-A7-Qi4 from the foundation block	219
73.2	Outward map: exports from GG-A7-Qi4 to downstream papers	221
74	The Ten Structural Innovations of the Qi-4 Runtime	222
74.1	Global naming convention for cross-cohort citation	224
75	Open Problems, Limitations, and Falsifiability	224
75.1	Why assessment must be split from significance	225
75.2	Open problems and limitations of the derivational stage	225
75.3	Falsifiability and empirical program	227
76	What a framework is entitled to claim	230
Part XII. Conclusions		232
77	Conclusions	232
77.1	The synthesis theorem	232
77.2	Three principal conclusions	234

77.3	Outlook and the next generation of papers	235
77.4	Closing reflections: four pillars, one object	236
Acknowledgments		241
Part XIII. Technical Appendices		242
A	Summary of Abbreviations	242
B	Glossary of Key Symbols	248
C	Key equations	249
D	Technical derivations for the GRIP pipeline	252
D.1	Derivation of the transformed generator	252
D.2	Quotient reduction of the continuity equation	252
D.3	Local bifurcation algebra	252
D.4	Catalog of possible Geometric Attractors	253
E	Technical derivations for Qi-QUA	254
E.1	Coefficient-space to natural-unit conversion	254
E.2	Numerical ledger of the Arisaka Constant	254
E.3	Coarse-grained cone generated by ideal locked events	255
E.4	Event-budget table	256
F	Technical derivations for Qi-EP	256
F.1	Reduction of the current-density action to a trajectory functional	256
F.2	Force-controlled extremum and the dual minimum-dissipation form	257
F.3	Pullback to the statistical manifold and the Fisher metric	259
F.4	A two-state blanket model	261
F.5	Dictionary of the three gauges	261
F.6	The duality theorem: both readings, in full	262
F.7	The failure theorem: the exact obstruction	264
F.8	The overdamped limit, and why it is legitimate	265
G	Technical derivations for Qi-CON	265
G.1	Radial integration and kernel homogeneity	265
G.2	Characteristic function and stable-law tail	266
G.3	Box-counting derivation of the fractal dimension	267
G.4	Spectral density and the two-knee approximation	268
G.5	Compact diagnostic dictionary	268
H	Technical derivations for Qi-HAL	269
H.1	Uniqueness of the MePMoS causal order	269
H.2	Metric recovery in NHT	269
H.3	Frequency-independent HAL readout	270
H.4	The dyadic ladder and level count of BFFT	270
H.5	Normal-form bifurcation for the DGI–GSSB switch	270
I	Technical details for SIP and GFF	271
I.1	Equivalence of folded and unfolded packet forms	271
I.2	Detailed derivation of the GFF balance law	271
I.3	Event-level reduction on the locked current ray	272
I.4	Proof sketch for the viability-tube criterion	272
I.5	Compact packet cards across representative domains	273
J	Technical derivations for OH, IH, and EIC	273

J.1	Derivation of the outer-hologram kernel from instruments	273
J.2	Observability and action-conditioned reconstruction	274
J.3	Two-view triangulation from lower-dimensional sensory inputs	275
J.4	Internal measurement operator and selected content	275
J.5	Lyapunov decrease of the EIC functional with bounded fluctuations	275
J.6	Representative OH/IH/EIC cards	276
K	Pagoda hierarchy and cross-domain technical details	277
K.1	Closure defects, effective floors, and error budgets	277
K.2	Hitting-time proof of conditional inevitability	277
K.3	Compact domain cards and observables	277
K.4	Minimal dynamical forms across domains	278
L	Falsifiability summary	279
L.1	Compact falsifiability matrix	279
M	Chemistry, biochemistry, and protocellular open systems	279
M.1	Why chemistry is the first nontrivial export layer	280
M.2	A chemistry-scale System Instance Packet	280
M.3	Reaction-network representation and current ledgers	281
M.4	Qi-EP in biochemical cycles	282
M.5	Qi-CON signatures in chemistry and protocells	283
M.6	From protocells to living systems	283
M.7	Concluding remark on the chemistry layer	284
N	Neural, cognitive, and consciousness-oriented open systems	285
N.1	A neural System Instance Packet	285
N.2	DGI, GSSB, and GA in neural dynamics	286
N.3	Qi-HAL and the minimal adaptive closure	286
N.4	Outer hologram, inner hologram, and evolving inner cosmos	287
N.5	Qi-CON and the spectrum of neural timescales	288
N.6	Scope, caution, and falsifiability	288
O	Ecological, social, and technological collective systems	290
O.1	A collective System Instance Packet	290
O.2	Dynamic equations and the meaning of the four channels	290
O.3	Qi-EP for collective adaptation	291
O.4	Qi-CON, heavy tails, and cascade structure	292
O.5	Where the framework can fail in collective domains	292
O.6	Why the collective appendix still matters	292
P	Notation concordance and benchmark protocols	293
P.1	Notation concordance adopted in the polishing round	293
P.2	Benchmark program	293
Part XIV. Common Objections and Responses		294
Q	For the curious reader: the big questions, plainly answered	294
R	The particle theorist and mathematical physicist	306
S	The systems biologist and biophysicist	310
T	The origin-of-life researcher	313
U	The neuroscientist	316

V	The methodologist and the working scientist	319
Part XV. References		321

Preface

In 1944, Erwin Schrödinger asked a question that physics has never quite answered: what is life? In 1948, Claude Shannon wrote down a quantitative measure of information that any engineer could compute. In 1961, Rolf Landauer showed that erasing a single bit of that information always costs at least a fixed amount of physical heat. In 2010, Karl Friston proposed that every living organism keeps itself alive by minimizing a free-energy functional, the way a thermostat minimizes a temperature difference. And since the work of 't Hooft in 1993 and Susskind in 1995, theoretical physicists have entertained the hypothesis that the world we see may be a projection of information stored on a distant boundary. Each of these threads was woven in a different decade, by people working in different fields, and the threads have never been tied together.

Living systems sit at the intersection of every one of these threads. A cell is a thermodynamic engine, an information-processing device, a bulk-to-boundary projection of a deeper geometry, and a free-energy minimizer, all at once. Physicists who study quantum mechanics speak of measurement records and decoherence; biologists who study cells speak of memory traces and homeostasis; theorists who study black holes speak of bits encoded on horizons. The vocabulary is utterly different, the underlying mathematics often quite similar, and the connections have usually been left to philosophical sidebars. There has been no common language in which to say that the information current in a cell and the information current on a holographic screen are two faces of one object.

This paper proposes a common language. The argument names four informational pillars on a small six-dimensional manifold: a master quantum-information object that supplies the bits, an equivalence principle that tells us a bit of information always costs the same minimum amount of energy no matter who is doing the bookkeeping, a master continuity statement that describes how information rolls forward in any open system, and a master hierarchy that explains why nature seems to climb the same ladder of organization, from molecules to memory, on every scale. These pillars do not add new postulates. They are the informational face of the geometric pillars derived in the companion paper of the program, and together with them they close the loop between geometry, gauge, energy, and information. A single conversion factor of size 0.1555 electron-volts per qubit connects the information current of a thinking brain to the energy current of the cell that hosts it — and matches one-half the energy released by the breakdown of a single molecule of ATP at body temperature (one ATP carries two such quanta).

We invite the reader to walk the four informational pillars in the same way the companion paper walks the four geometric ones. Read each pillar, follow each proof, and ask whether the laboratory phenomenon it predicts — the conversion factor between information and energy, the spectral signature of memory, the universality of the ladder — is something a careful experiment could test. After all, the question Schrödinger asked in 1944 was whether life requires a new kind of physics. The answer that emerges here is that life does not require new physics — but it does require physics to recognize that information has always been one of its first-class objects, and that the boundary on which we live is, in a precise and laboratory-testable sense, the place where the upstream geometry writes itself down in bits.

Part I

Entrance and bridge

Introduction, recap, and the Qi-runtime extension claim

This Part answers one question: why should a single law describe a particle detector, a bacterium, a nervous system and a market?

The usual answer is that it should not, and the instinct behind that is a good one. Those four things differ in every way that normally matters. What this Part argues is that they share the one property that turns out to be decisive, and that the property has nothing to do with how complicated they are. Each is studied by someone who can only see part of it. A finite observer, using finite instruments, never has the whole state of anything in hand — and the moment that is true, the description available to that observer cannot be closed. It needs memory, or hidden variables, or averaging. The result is proved here rather than assumed, and it has a blunt consequence: the difference between physics and biology is not that one is closed and the other open. Both are open. Only the richness of what is left after the projection differs.

That is the entrance. The bridge is the other half of the Part. This paper derives nothing about the six-dimensional geometry, nothing about the boundary law, and nothing about the engine that drives a system downhill; all of that arrives from five companion papers, and this Part sets out exactly what arrives, in the form it will be used, with the theorem that supplies each piece named. That inventory is written as a contract rather than as a summary, because a later Part that quietly assumed something not on the list would be borrowing without saying so. One item on the list carries a condition of its own — the runtime cascade closes subject to two conjectures still open in its home paper — and everything built on it in this paper inherits that condition.

The Part then does one more thing, and it is the reason the paper exists. It shows that the boundary law, on its own, is not enough. A law can describe how an open system evolves and still leave the words we use about such systems — memory, selection, persistence, complexity — attached to nothing in particular, so that the same system admits many incompatible readings. What is missing is a level between the law and the domain: a runtime. The rest of the paper builds it.

A reader in a hurry should not skip the table of claims and non-claims. It is short, it is the least exciting page in the Part, and it is the one that says what this framework does not assert. Vocabulary this broad is easy to over-read, and the right-hand column of that table is what keeps the left-hand column sayable.

Part I sets the stage for everything that follows. We open with the question that physics has never quite answered: what is life, and more generally, what is an open system? Section 1 states the problem in plain language and lists the unsolved questions that motivate the present paper; Section 3 recapitulates the upstream foundation (Papers I–IV-A) on which the Qi-4 runtime must be built; and Section 4 states the universal-extension claim — that GDOS applies to every open

system in nature once the four Qi-4 pillars are in place. Part I tells the reader why the rest of the paper is worth reading.

1 Introduction

1.1 Why living organisms are the natural entry point for Paper IV-B

The first four papers of the GG/GD sequence were necessarily written from the bottom up. Paper I fixed the constitutional layer: what a lawful theory is allowed to assume, what it must derive, and what it must make measurable (Arisaka, 2026, A1-Constitution). Paper II then answered the ultraviolet question, namely what higher-dimensional geometry is forced by that constitution (Arisaka, 2026, A2-GG6). Paper III answered the operational question, namely how a finite observer on the boundary obtains quantum states, measurement records, and four-dimensional effective dynamics from that higher-dimensional bulk (Arisaka, 2026, A5-GDOS). Paper IV-B must now reverse the direction of emphasis. It does not begin from the smallest scales but from the systems in nature that most vividly display the full open-system problem: living organisms.

A living organism is the most stringent natural test case for any theory that claims to be a universal theory of open systems. It is not merely a dynamical system, nor merely a thermodynamic system, nor merely an information-processing system. It is all three at once. A cell or organism maintains a boundary, exchanges matter and energy with its environment, processes information, stores memory, predicts possible futures, selects actions, and persists far from equilibrium for times vastly longer than microscopic relaxation times (Friston, 2010, Landauer, 1961, Nicolis and Prigogine, 1977, Prigogine, 1967, Schrödinger, 1944). If a framework cannot say what such a system *is*, what its lawful state packet is, what its relevant currents are, why it persists, and how it makes or receives measurements, then that framework has not yet reached a universal runtime.

This is why life, rather than elementary particles alone, is the natural entry point for the present paper. Particles reveal the constitutional grammar of the theory. Organisms reveal whether that grammar is strong enough to handle persistence, adaptation, memory, agency, and internal modeling without changing its foundations from one scientific domain to the next. In that sense the present paper is not a departure from Papers I–IV-A but the first large-scale stress test of their universality.

Mathematically, the key reason organisms are unavoidable is that they make all hidden assumptions visible. To make this concrete, we now write the minimal quantum-information notation that a finite observer is forced to use, even when modeling a cell or an organism. A reader from biology or chemistry who is unfamiliar with this notation may skim the equation and continue with the prose immediately after it; the gloss following the equation translates each symbol into plain words.

For an isolated Hamiltonian system the notation

$$\frac{d}{dt}|\psi(t)\rangle = -iH|\psi(t)\rangle$$

can be written compactly and elegantly, but this compactness hides the fact that a real observer never has access to the entire Hilbert space. In an actual experiment or an actual organism, one deals instead with a subsystem S , an environment E , a reduced state, and a family of coarse-grained

observables:

$$\mathcal{H}_{\text{Bulk}} = \mathcal{H}_{\text{sys}} \otimes \mathcal{H}_{\text{env}}, \quad \rho_S(t) = \text{Tr}_E \rho_{SE}(t), \quad \langle O_a \rangle_t = \text{Tr}_{\mathcal{H}_{\text{sys}}} (\rho_S(t) O_a). \quad (1.1)$$

In plain words: $\mathcal{H}_{\text{Bulk}}$ is the full state space of everything — system plus environment together; $\mathcal{H}_{\text{sys}}$ is the part the observer can access; $\mathcal{H}_{\text{env}}$ is everything else. The reduced state $\rho_S(t)$ is what the observer would measure if she traced over (i.e. averaged over) the environment she cannot resolve. The expectation value $\langle O_a \rangle_t$ is the actual number a laboratory experiment returns for the observable O_a at time t , computed from the reduced state. None of this notation is optional once one accepts that the observer is finite; it is the precise bookkeeping that the rest of the paper will work with.

A living organism makes Eq. (1.1) unavoidable because it is *constitutively coupled* to an environment: food, photons, chemical gradients, thermal baths, predators, conspecifics, and internal submodules all act as ports. The closed-system idealization is therefore not just inaccurate but structurally misleading when one attempts to understand life, cognition, and collective behavior.

The philosophical point is equally important. Since [Schrödinger \(1944\)](#), life has been recognized as a phenomenon that cannot be reduced to static equilibrium thermodynamics. Since [Shannon \(1948\)](#) and [Landauer \(1961\)](#), information can no longer be treated as an informal metaphor detached from physical cost. Since the development of nonequilibrium thermodynamics and open quantum systems, irreversibility, coarse-graining, and environment-induced records must be taken as first-class theoretical objects rather than nuisances to be postponed ([Breuer and Petruccione, 2002](#), [Nicolis and Prigogine, 1977](#), [Prigogine, 1967](#)). Paper IV-B begins from organisms because they force all these strands to coexist in one object.

Table 1: Why living organisms are the natural entry point for the universal extension of GDOS. Five features of living matter are listed down the first column; the second says why each is mathematically demanding, and the third why each is decisive for the present paper. The argument the table carries is not that biology is special but that it is *unavoidable*: a system that is persistently open, held far from equilibrium, maintains its own boundary, remembers and predicts, and is organized hierarchically exercises every part of the runtime at once, so it is the case in which a missing piece of the framework shows up first.

Feature of living systems	Why it is mathematically demanding	Why it is decisive for Paper IV-B
Persistent openness	State evolution requires explicit environmental coupling, ports, and reduced dynamics.	Forces the theory to move beyond closed-system Hamiltonian language.
Far-from-equilibrium operation	Stability is maintained by continuous flux, not by equilibrium minimization alone.	Motivates a unified free-energy ledger and trajectory-selection principle.
Boundary maintenance	The system must distinguish self from environment while exchanging with it.	Motivates SIP, SeOS, and the outer/inner hologram distinction.
Memory and prediction	The present state depends on history and on internally represented futures.	Motivates Qi-HAL and the move from external measurement to internal measurement.
Hierarchical organization	Molecular, cellular, organismal, and social levels interact without collapsing into one scale.	Motivates scale covariance, Qi-CON, and the pagoda hierarchy.
Action and adaptation	The system does not only receive information; it changes its boundary conditions by acting.	Motivates the full DGI→GSSB→GA pipeline rather than passive state evolution alone.

1.2 From the measurement problem to the universal observation problem

Paper III was framed around the measurement problem in quantum mechanics, but its deeper lesson is broader. The familiar measurement problem asks why and how definite outcomes appear if the microscopic law is unitary (Bell, 1964, von Neumann, 1932, Zurek, 2003). The answer proposed in Paper III was to promote measurement from an interpretational add-on to a constitutional demand: AX-3 (OGGEP) requires every admissible theory to provide a lawful bulk-to-boundary dictionary (Arisaka, 2026, A1-Constitution; Arisaka, 2026, A5-GDOS). In its most compact form, this dictionary is a map

$$\mathcal{D}_{\text{OGGEP}} : \text{bulk data} \longrightarrow (\mathcal{H}_{\partial}, \rho_{\partial}, \Phi_{t,s}, \{\mathcal{I}_{\alpha}\}), \quad (1.2)$$

where ρ_{∂} is the observer-accessible boundary state, $\Phi_{t,s}$ is an admissible evolution map, and $\mathcal{I}_{\alpha}$ are instruments or operations associated with registered outcomes.

For Paper IV-B the crucial point is that Eq. (1.2) is not merely a statement about laboratory measurements in quantum physics. It is the general grammar of *observation*. Every scientific observation, in any field, is obtained only after three reductions have already occurred:

- (i) a large inaccessible state space has been truncated to a smaller accessible state space;
- (ii) a continuous high-dimensional process has been projected onto a finite register or data stream;

(iii) the resulting record has been interpreted relative to a chosen instrument class.

That is as true for electrophysiology, microscopy, ecology, and economic indicators as it is for Stern–Gerlach experiments. The observation problem is therefore universal. The quantum measurement problem is only its sharpest historically visible corner. The openness, however, is not created by the act of study: the bulk-to-boundary projection of OGEP is always on (GG-A5-GDOS §23), running at every boundary point whether or not anyone is studying anything, and a finite observer merely *inherits* an already-projected, already-open description. What is universal is not that observation makes systems open, but that every system *is* open — observed or not.

The mathematical signature of universality is that the same reduced-dynamics structure recurs across domains. At the exact level one obtains a non-Markovian memory-kernel equation of Nakajima–Zwanzig type,

$$\frac{d}{dt}\rho_{\partial}(t) = \int_0^t K(t-s)\rho_{\partial}(s)ds + I(t), \quad (1.3)$$

while in the short-memory limit one obtains a completely positive semigroup with GKLS generator,

$$\frac{d}{dt}\rho_{\partial} = -i[H_{\text{eff}}, \rho_{\partial}] + \sum_k \left(L_k \rho_{\partial} L_k^{\dagger} - \frac{1}{2} \{L_k^{\dagger} L_k, \rho_{\partial}\} \right). \quad (1.4)$$

The exact microscopic meaning of H_{eff} , L_k , and the kernel K varies across domains, but the open-system logic does not. What changes from particles to organisms is not the fact of openness but the complexity of the state packet, the boundary, the port structure, the attractor geometry, and the internal modeling architecture.

This is precisely why the present paper must broaden the focus from “measurement” to “observation.” An organism not only registers external data; it also constructs internal states that function as measurements of its own world-model. A nervous system, for example, receives sensory inputs, compares them against memory, generates predictions, issues actions, and re-samples the environment. Once that loop is recognized, external observation and internal observation become two instances of the same architecture. This is the conceptual doorway through which OH+IH = EIC will later enter: the outer hologram concerns externally registered boundary data, while the inner hologram concerns the internally reconstructed model through which a system experiences and navigates its world.

Definition 1.1 (Operational openness). A scientific system S is *operationally open* if the observer-accessible state of S cannot be described by a closed autonomous law on S alone, but instead requires at least one of the following: a traced-over environment, explicit ports, memory kernels, stochastic coarse-graining, or instrument-conditioned state updates.

Remark 1.2. Operational openness is weaker than full ontological openness. The definition makes no claim that the total universe is open to something external. It claims only that every *scientific description used by a finite observer* is open because it is necessarily partial, coarse-grained, and instrument-bounded. This is enough to make open-system mathematics constitutional.

1.3 The need for a common runtime

The notion of a *runtime* is meant literally. Papers I–IV-A supplied the constitution, the ultraviolet hardware, and the boundary operating system. Paper IV-B now asks: what is the executable layer common to all natural open systems once the theory is actually running? In software language, this is the layer that every scientific domain calls even if the domains look superficially different. In physical language, it is the layer that every open system instantiates even if the microscopic details differ.

The proposed answer is the master pipeline

$$\boxed{\text{DGI} \longrightarrow \text{GSSB} \longrightarrow \text{GA}} \quad (1.5)$$

together with the Qi-4 pillars and the universal runtime vocabulary. Equation (1.5) should be read as the minimal grammar of persistence. First, a system explores lawful gauge-covariant possibilities (DGI). Second, environmental coupling, constraints, and boundary conditions select one branch or one structure (GSSB). Third, the selected structure is stabilized as an attractor, function, phenotype, memory, or policy (GA). The later sections of the paper make each arrow explicit in the fixed order adopted here. For the introduction, it is enough to see that persistence is not a static property. It is a three-stage process.

The Qi-4 pillars then answer four questions about any runtime instance:

Qi-QUA *What currents are locked together?*

Qi-EP *Which trajectories are selected?*

Qi-CON *What scale signatures should appear?*

Qi-HAL *What architecture carries memory, prediction, sensing, and action?*

This decomposition is not ornamental. It is a discipline against conceptual drift. Without it, one risks mixing bookkeeping, variational principles, statistical diagnostics, and architecture into one undefined notion of “complexity.”

A convenient preview of the mathematical structure is the generic GDOS continuity form already derived in Paper III:

$$\partial_t p(x, t) = -\nabla_A J^A(x, t), \quad (1.6)$$

with current decomposed schematically as

$$J^A = -D^{AB}\nabla_B p - \mu^{AB}p\nabla_B \mathcal{F}_{\text{GFF}} + \Omega^A p. \quad (1.7)$$

The first term represents diffusive spreading, the second gradient descent on an effective GFF ledger, and the third divergence-free circulation. The logic of Paper IV-B is that a universal open-system runtime cannot be built from any one of these terms alone. Persistence needs all three. Pure diffusion destroys structure; pure gradient descent over-freezes structure; pure circulation never commits structure. The master pipeline $\text{DGI} \rightarrow \text{GSSB} \rightarrow \text{GA}$ is the conceptual version of this mathematical trichotomy.

Definition 1.3 (Runtime packet preview). A *runtime packet* is a preliminary placeholder for what will later be formalized as a **SIP**. It is the minimal declared data set required to specify a scientific open system:

$$\mathcal{R} = (\mathcal{X}, \mathcal{P}, \mathcal{O}, \rho, \Phi, \mathcal{J}, \mathcal{F}, \mathcal{A}_{\text{arch}}), \quad (1.8)$$

where $\mathcal{X}$ is the state manifold, $\mathcal{P}$ the environmental interfaces, $\mathcal{O}$ the registered observables, ρ the effective state, Φ the effective evolution law, $\mathcal{J}$ the current ledger, $\mathcal{F}$ the scalar selection functional, and $\mathcal{A}_{\text{arch}}$ the architecture responsible for memory, prediction, sensing, and action.

The full paper will refine Eq. (1.8) substantially. The point here is merely to show why a common runtime is needed: without some declared packet of this kind, the same words—“state,” “information,” “stability,” “selection,” “complexity”—mean incompatible things in different domains.

1.4 Preview of the new objects introduced in Paper IV-B

It is useful to state at the outset what is genuinely new in Paper IV-B relative to Paper III. Paper III derived the existence of a lawful open boundary theory. Paper IV-B extends that theory by adding the runtime objects needed for universality. These objects are not optional decorations; they are the minimal extensions needed to move from the foundational level of measurement and particle physics to the domain-general level of chemistry, life, mind, and society.

First, Paper IV-B introduces **DGI**, **GSSB**, and **GA** as the universal dynamical spine. Although their ingredients are already latent in Papers I–IV-A, they were not yet elevated to a named, reusable runtime grammar. Paper IV-B makes this promotion explicit.

Second, Paper IV-B introduces the **Qi-4** pillars as a disciplined factorization of the runtime. **Qi-QUA** is the bookkeeping layer; **Qi-EP** is the selection layer; **Qi-CON** is the diagnostic layer; **Qi-HAL** is the architecture layer. This factorization is what allows one to move from particle measurements to molecular machines and then further to organisms and brains without rewriting the language each time.

Third, Paper IV-B introduces **SIP**, **GFF**, and **SeOS** as the working vocabulary of universal open-system analysis. A **SIP** will later serve as the minimal packet that can be reused across scientific branches. **GFF** will later serve as the combined free-energy ledger. **SeOS** will later designate persistent self-maintaining open systems whose continued existence is itself part of the dynamical story.

Fourth, Paper IV-B introduces **OH + IH = EIC**, which extends the meaning of measurement. In the first three papers, the concern was primarily with how a physical observer obtains a lawful boundary record. In the fourth paper, the concern broadens to how a living or cognitive system constructs an internal world-model and how that world-model contributes to what is subjectively experienced.

Finally, Paper IV-B introduces a hierarchy picture. The Japanese five-story pagoda is not used merely as a poetic image. It is intended as a compact structural map: the same **DGI** spine runs through different layers of **GSSB** output, from particles to cosmos to chemistry to life to mind. This hierarchical picture is important because a universal theory of open systems must explain not only one level but also why multiple levels can coexist without collapsing into each other.

1.5 Claims and non-claims of the present paper

The main positive claim is that there exists a single geometric runtime capable of organizing open systems across domains, once the groundwork of Papers I–IV-A is accepted. The paper does *not* claim that all sciences reduce to particle physics, that every empirical detail of biology/neuroscience/society is already solved, or that earlier frameworks are useless. It claims only that the admissibility grammar—how states, currents, observations, attractors, and hierarchies are declared and audited—can be unified, with geometry and admissibility first, runtime second, and domain-specific models third. Table 2 fixes the contrast.

Table 2: Claims and non-claims that discipline the interpretation of this paper. Five categories run down the first column — universality, reduction, measurement, life and mind, and the relation to prior theories — with what is claimed in the second and what is explicitly *not* claimed in the third. The right-hand column is the load-bearing one. It is placed this early because the framework’s vocabulary is broad enough to be over-read, and a reader who takes the non-claims with them will not mistake an organizing statement for a predictive one.

Category	Claim of Paper IV-B	Explicit non-claim
Universality	One common runtime can organize open systems across domains.	Not every domain-specific equation is already fully derived in this step.
Reduction	A common admissibility grammar exists from particles to society.	The paper does not reduce away all higher-level descriptions.
Measurement	Observation is treated as a lawful open-system projection.	The paper does not identify measurement with consciousness.
Life and mind	Life and cognition are treated as lawful runtime outcomes.	The paper does not claim that all biological and cognitive details are already solved.
Relation to prior theories	Existing frameworks are reinterpreted as partial slices of a larger runtime.	The paper does not claim that their local successes are invalid.

1.6 Roadmap of the full manuscript

The full paper is designed as a landmark synthesis rather than a short note. Its architecture is a fifteen-Part structure: twelve Parts of main argument followed by a three-Part back matter. Part I supplies the entrance and bridge: the Introduction, the recap of Papers I–IV-A, the place of the paper in the twenty-two-paper cohort, and the universal extension of GDOS to any open system. Part II imports the runtime constitution from GG-A6-GRIP in the strict order demanded by logical dependence: first the master theorem $\text{GRIP} = \text{DGI} \rightarrow \text{GSSB} \rightarrow \text{GA}$, then the scalar ledger GFF, and only then the canonical system-instance packet SIP. Six Parts then develop the Qi-4 runtime, one subject each, in its own dependency order: current locking (Part III), universal extremality (Part IV), multiscale diagnostics (Part V), and memory-bearing architecture (Part VI), followed by the persistence and completeness layer (Part VII) and the statement tying the four pillars to the

ground attractor of the geometric companion (Part VIII). Part IX builds the internality layer: the Outer and Inner Holograms and the internal-measurement runtime ($OH + IH = EIC$). Part X treats emergence, the five-story pagoda, the cross-domain applications, and the explicit universality claim. Part XI is the assessment: significance and comparison, limitations, and falsifiability. Part XII draws the conclusions. The back matter then follows in three further Parts: the Technical Appendices (Part XIII), the Common Objections and Responses (Part XIV), and the References (Part XV).

The present integrated manuscript already executes this entire order in the main text. It should therefore be read not as a collection of separately drafted chapters but as one continuous derivational arc: bridge $\rightarrow$ runtime $\rightarrow$ scalar budget $\rightarrow$ packet $\rightarrow$ Qi event locking $\rightarrow$ selector $\rightarrow$ diagnostics $\rightarrow$ architecture $\rightarrow$ persistence/completeness $\rightarrow$ internality $\rightarrow$ hierarchy $\rightarrow$ applications $\rightarrow$ universality $\rightarrow$ assessment. The purpose of the present harmonization pass is to make every transition, theorem lead-in, and section summary reflect that non-interchangeable order as clearly as possible.

2 This paper’s place in the cohort: one link of the open-system chain

This paper establishes the four quantum-information Qi-4 pillars and the operational meaning of measurement; that is its own subject. For readers following the whole GG-Theory cohort, it is also one link of a single open-system chain that the consequence papers cash out. Because no physical system is closed, the four-dimensional world is a lawful projection (OGGEP) of the six-dimensional bulk, and that projection is observation itself; its dynamics is the boundary law GDOS, forced by the geometric Gi-4 and quantum-information Qi-4 pillars, so that the laws, the particles, and the contents of the universe are *observables* of the projection rather than free inputs. In that chain, the present paper supplies the Qi-4 pillars and the measurement content — the part of the chain that makes the bulk-to-boundary projection *observation*. The funnel below is shared, in structure, with the companion papers GG-A9-SM, GG-A10-Cosmos, and GG-A11-Origin; here it is offered only as orientation, and the rest of this paper proceeds on its own terms.

3 Recap bridge to GG-A1-Constitution, GG-A2-GG6, GG-A3-BI6, GG-A5-GDOS, and GG-A6-GRIP

3.1 Why the recap is structurally necessary

Paper IV-B is not intended as an isolated monograph. It is the fourth member of a logically ordered sequence. The whole point of that sequence is that later papers should not re-open earlier constitutional questions every time a new domain is addressed. If that discipline is not maintained, the program would degenerate into a loose collection of ambitions rather than a theorem chain.

The recap therefore plays a very specific role. It does not repeat the earlier papers for convenience. It identifies which results from those papers are now to be treated as *imported theorems*, which notation choices are fixed, which issues were intentionally postponed, and which exact gap makes




Figure 1: The open-system chain shared across the GG-Theory cohort. Because no physical system is closed, the four-dimensional world is a lawful projection (OGGEP, AX-3) of the six-dimensional bulk, and that projection is observation itself; its dynamics is GDOS, forced by the Gi-4 and Qi-4 pillars, so the laws and contents of the universe are observables rather than free inputs. The same funnel opens the consequence papers GG-A9-SM, GG-A10-Cosmos, and GG-A11-Origin, with a paper-specific payload; within it the present paper establishes the Qi-4 pillars and the measurement content.

Paper IV-B necessary. In that sense the recap is a bridge, not a summary. It is the interface contract between the already established foundation and the new runtime layer.

A concise way to say this is that each paper answers a different question:

- Paper I : What counts as an admissible theory?
 - Paper II : What ultraviolet bulk is forced by that admissibility?
 - Paper III : How does a lawful boundary world arise from that bulk?
 - Paper IV : How do persistent open systems run within that lawful boundary world?
- (3.1)

Equation (3.1) already shows why Paper IV-B cannot merely be an appendix to Paper III. The question it asks is different in kind. Papers I–IV-A establish legality, stage, and projection. Paper IV-B establishes runtime.

Table 3: Division of labor across the papers this one stands on. Each row is one companion: the second column gives the primary question that paper answers, the third its principal theorem-level output, and the fourth the direct role that output plays here. The table exists to make the direction of the debt explicit — everything in the first three columns is imported and none of it is re-derived here — and it is the reference a reader should use when a later section says a result is inherited.

Paper	Primary question	Principal theorem-level output	Direct role in Paper IV-B
GG-A1-Constitution	What is a lawful theory?	Four axioms, seven postulates, admissibility grammar	Fixes the constitution and the ordered filters Paper IV-B must obey.
GG-A2-GG6	What ultraviolet bulk is forced?	6D bulk, BI-6, OGN-5, anomaly-free 6D QFT	Fixes the geometric stage, current ordering, and ultraviolet inputs.
GG-A5-GDOS	How does the boundary world arise?	OGGEP channel, GDOS, 4D QM, 4D GR, formal IR QFT/SM descent	Fixes the open-system boundary law and the meaning of measurement.
A6/A7 (runtime-constitution program)	How do open systems persist and organize?	Runtime grammar, Qi-4 pillars, SIP/GFF/SeOS, OH/IH/EIC hierarchy	Extends the boundary law into a universal theory of open systems.

3.2 Paper I recap: the **GG-Constitution** axiom system

Paper I established that the program is to be judged first as a constitution rather than as a model (Arisaka, 2026, A1-Constitution). A candidate theory $\mathcal{T}$ is admissible only if $\text{Admissible}(\mathcal{T}) := (\mathcal{T} \models \text{AX}) \wedge (\mathcal{T} \models \text{POS})$, with the four axioms AX-1 (LC), AX-2 (GGEP), AX-3 (OGGEP), AX-4 (MSMF) fixing primitive semantics, and the seven postulates POS-1–POS-7 enforcing gauge invariance, a unified extremal principle, anomaly freedom, closure/positivity, scale covariance, OGN, and measurement/falsifiability respectively—in that order. For Paper IV-B this matters not as a slogan but because the constitution is imported as Theorem T-AX of Paper I; see also the inward Coherence Map (Sec. 73) for the line-by-line audit.

3.3 Paper II recap: **GG-6**, **BI-6**, and the six-dimensional bulk

Paper II solved the ultraviolet realization problem (Arisaka, 2026, A2-GG6), deriving a minimal six-dimensional bulk

$$M_6 = M_4 \times I_r \times S_\theta^1, \quad r = \ln(E_{\text{GUT}}/\mu), \quad ds_6^2 = e^{-2A(r)} g_{\mu\nu} dx^\mu dx^\nu + dr^2 + R_\theta^2 d\theta^2, \quad (3.2)$$

with bulk gauge structure $G_{\text{Bulk}} = \text{SU}(5) \times \text{U}(1)_B \times \text{U}(1)_{Q_i} \times \text{Gravity}$. The technical centerpiece is BI-6, $dH_3 = \text{tr}(R \wedge R) - \text{Tr}(F_5 \wedge F_5) - (2/5)F_B \wedge F_B - (2/5)F_{Q_i} \wedge F_{Q_i}$, with coefficient vector $(a, b, c, d) = (1, -1, -2/5, -2/5)$ and the OGN-5 integer ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$ satisfying $K = LJ = M^2 + N^2 + L^2 = 35$. These are imported as Theorem T-BI6 of Paper II; see the inward Coherence Map. The four-component bulk is constitutionally admissible in A1–A3 but *operationally required* for Qi-4: the Arisaka equation contains J_{Q_i} explicitly, Qi-HAL treats $\text{U}(1)_{Q_i}$ as the memory phase carrier, and the entire OH+IH=EIC framework is defined on memory-bearing

$U(1)_{Q_i}$ systems (a three-component variant would require a different Qi-4' framework, recorded as open problem O9). The $U(1)_{Q_i}$ circle is a real compact fiber; the phase it carries is a real angle, and the imaginary unit of that phase's boundary representation is the derived quarter-turn of T-P4-1 (Arisaka, 2026, P4-Real) — so Qi-HAL's memory register is a real angle on a real circle, its complex dressing inherited rather than assumed.

3.4 Paper III recap: GDOS as the boundary foundation

Paper III answered the projection problem (Arisaka, 2026, A5-GDOS). With $\mathcal{H}_{\text{Bulk}} = \mathcal{H}_{\partial} \otimes \mathcal{H}_{\text{env}}$ and $\rho_{\partial}(t) = \text{Tr}_{\mathcal{H}_{\text{env}}} \rho_{\text{Bulk}}(t)$, the exact projected dynamics obeys a Nakajima–Zwanzig memory-kernel equation that reduces in the Markovian limit to a GKLS generator and, in the mesoscopic coarse-grained limit, to a continuity law

$$\partial_t p = -\nabla_A J^A, \quad J^A = -D^{AB} \nabla_B p - \mu^{AB} p \nabla_B \mathcal{F} + \Omega^A p. \quad (3.3)$$

Equation (3.3) is the kernel from which Paper IV-B grows. Paper III recovered 4D quantum mechanics (with measurement as a lawful instrument update), 4D general relativity (from the metric zero mode), formal IR QFT, and an IR Standard-Model sector with $SU(3) \times SU(2) \times U(1)_{EM} \times U(1)_B \times U(1)_{Q_i}$ structure ($U(1)_B$ enforcing exact proton stability). Paper III was intentionally foundational: it did not yet build the universal runtime needed for chemistry, life, brains, or society.

3.5 Paper IV-A recap: GRIP, GFF, and SIP

Paper IV-A made the runtime constitution explicit (Arisaka, 2026, A6-GRIP): the Gauge-Runtime Invariance Pipeline $\text{GRIP} := (\text{DGI} \rightarrow \text{GSSB} \rightarrow \text{GA})$, the derived scalar ledger $F_{\text{GFF}} = F_{\text{GEP}} + F_{\text{CEP}}$ (with GEP descending from the geometry–gauge sector of the 6D action and CEP descending from boundary state-manifold curvature/diffusion), and the canonical system-instance packet **SIP**. The ledger GFF and the pipeline GRIP are imported as GG-A6-GRIP's principal theorems T3 and T4 respectively (its T1 and T2, TGG and CCI, supplying the Master Conservation Law and the OGN-5 ledger), while SIP is imported as GG-A6-GRIP's system-instance-packet construction rather than as a principal theorem. As established in GG-A6-GRIP, the GRIP-cascade closure (T4 — including the $N_g \leq 3$ statement and the dual-layer compatibility extension $\mathfrak{C}_{\text{GRIP}}$) holds subject to GG-A6-GRIP's cascade conditioning: the compensator monodromy statement M11, carried there in spectral-flow form with a verified benchmark instance — M10 itself, the BI-6 budget bound, is resolved externally by the capacity theorem T-N3-6 of Arisaka (2026, P2-N3), with saturation an arithmetic identity, the topological residue discharged on the benchmark, and the scope refinements closed at class level by that companion's rigidity theorems; the constructions of the present paper that rest on GRIP inherit that conditional status. The ground-attractor existence and uniqueness result is imported separately and does not depend on any cascade residue: it is GG-A6-GRIP's per-stratum ground-state corollary, which holds on the strata that paper calls admissible — compact, connected, Kähler, with the free-energy functional restricted to the stratum the moment map of a Hamiltonian torus action with isolated fixed points — and the last of those clauses is a physical identification, graded as one by that paper and derived there on the family stratum only. That grade travels with every use

made of the result here. No Qi-runtime can be introduced before GRIP, GFF, and SIP are in hand; the present paper begins from this imported runtime constitution and asks what new theorem layers become inevitable once informational current, persistent memory, and internal reconstruction are allowed to operate on top of it.

3.6 Bridge equation to Paper IV-B

The logical bridge from Papers I–IV-A to the present paper is

$$\boxed{(\text{GG-Constitution}) \implies (\text{GG-6 bulk with BI-6 and OGN-5}) \xrightarrow{\text{OGGEP}} (\text{GDOS boundary law}) \implies (\text{runtime of open systems})} \quad (3.4)$$

whose target is the set $\{\text{DGI, GSSB, GA, GFF, SIP, Qi-QUA, Qi-EP, Qi-CON, Qi-HAL, SeOS, QCI, OH, IH, EIC}\}$. The updated Arisaka current notation is locked once and used consistently throughout:

$$\text{ArisakaRatio} : J'_E : J'_e : J'_B : J'_{Qi} = 1 : -1 : -\frac{2}{5} : -\frac{2}{5} = 5 : -5 : -2 : -2, \quad (3.5)$$

$$\text{ArisakaEquation} : \frac{J_E}{Q_B} = J_e = J_B = J_Q, \quad J_Q \equiv J_{Qi}, \quad (3.6)$$

$$\text{ArisakaConstant} : Q_B = \frac{2}{5} \frac{E_{\text{Ryd}}}{K}, \quad K = 35. \quad (3.7)$$

The full derivations and interpretations appear in Part III.

Table 4: Consolidated imports into this paper from the foundation block. The first column names the imported object, the second the companion paper and theorem that establishes it, and the third what it is used for here. Read as a checklist it defines the paper’s dependency surface: if one of these objects fails upstream, the rows that use it fail here, and the third column says which. Nothing outside this table is assumed without being derived.

Imported object	Source	Role in GG-A7-Qi4
Admissibility filter (AX-1–4 + POS-1–7)	GG-A1-Constitution, Theorem T-AX	Constrains every runtime layer (DGI, GSSB, GA, GFF, SIP, Qi-4, SeOS/QCI, OH/IH/EIC).
6D bulk + G_{Bulk}	GG-A2-GG6, Eq. (3.2)	Supplies the lawful UV stage projected by OGGEP.
BI-6 coefficient (1, −1, −2/5, −2/5)	GG-A2-GG6 / GG-A3-BI6, Theorem T-BI6	Ultraviolet source of the locked Arisaka Ratio Eq. (3.5).
OGN-5 ledger (1, 3, 5; 7, 35)	GG-A2-GG6 / GG-A3-BI6	Prevents runtime objects from reintroducing continuous freedom; supplies $K = 35$ for Q_B .
GDOS boundary law	GG-A5-GDOS, Eq. (3.3)	Kernel for the universal runtime extension of Sec. 4.
GRIP, GFF, SIP	GG-A6-GRIP: GFF = T3, GRIP = T4 (T4 closure conditional on M11, spectral-flow form; M10 resolved externally, scope refinements closed at class level); SIP construction	Runtime constitution (Part II); not re-derived in the present paper.

Proposition 3.1 (The post-Paper III runtime gap). *Let $\mathfrak{F}_3$ denote the full theorem-level content of GG-A5-GDOS. Then $\mathfrak{F}_3$ is sufficient to define lawful open-system boundary evolution and measurement, but it is not sufficient by itself to classify, compare, and derive persistent complex open systems across chemistry, life, mind, and society.*

Proof. A reduced boundary law fixes the evolution of p_t but is invariant under arbitrary reparametrizations of the basin label, history record, and latent informational variable. Bistable open systems with multiple metastable states and history-dependent future dynamics therefore admit infinitely many “runtime” interpretations (current ledger, packet identity, branch-selection semantics, memory architecture) at the level of $\mathfrak{F}_3$ alone. The runtime layer of Paper IV-B—Qi-QUA, SIP, GRIP, Qi-HAL, etc.—supplies exactly the missing structure that resolves this non-uniqueness. ■

4 Universal extension of GDOS to any open system in nature

4.1 From the boundary law of Paper III to a universal runtime

Paper III derived GDOS in the form required for quantum measurement, four-dimensional field theory, and the infrared limit of particle physics (Arizaka, 2026, A5-GDOS). The resulting equations were written either as exact reduced channels, as GKLS generators, or as coarse-grained continuity equations on an effective state manifold. In all three representations the crucial structural fact was the same: once the inaccessible variables are traced over, the observer-accessible system is governed by an open dynamics rather than by a closed autonomous Hamiltonian law.

Paper IV-B now asks whether that structure was specific to physics laboratories or whether it was already more general than it first appeared. The answer proposed here is that the physics derivation was already universal in form. What changes from one branch of science to another is not the existence of an open-system law, but only the meaning of the accessible coordinates, the port structure, the relevant currents, the admissible observables, and the characteristic attractors. The task of the present section is therefore to strip the derivation down to the reusable kernel that survives across domains.

Let the exact boundary-visible state at the end of the OGGE reduction be denoted $\rho_\partial(t) \in \mathcal{D}(\mathcal{H}_\partial)$. Any finite scientific description then selects only a finite family of observables O^A , $A = 1, \dots, N$, and therefore only a finite family of coordinates

$$x^A(t) := \Gamma^A[\rho_\partial(t)] \equiv \text{Tr}(\rho_\partial(t) O^A), \quad A = 1, \dots, N. \quad (4.1)$$

Equation (4.1) is the true universal entrance point. In quantum physics the O^A may be projectors, pointer observables, or low moments of fields. In chemistry they may be concentrations, fluxes, or bond-state populations. In cell biology they may be copy numbers, membrane potentials, or methylation states. In neuroscience they may be firing rates, phase variables, or synaptic efficacies. In social systems they may be densities, prices, communication rates, or network-order parameters. The common feature is that the accessible description is always a coarse observable dictionary (Breuer and Petruccione, 2002, van Kampen, 2007).

Once Eq. (4.1) is written, the exact operator equation of Paper III induces a reduced evolution on an effective state manifold $\mathcal{X}$ with local coordinates x^A . The resulting probability density $p(x, t)$ is defined so that for every smooth test function f ,

$$\int_{\mathcal{X}} f(x) p(x, t) d\mu(x) = \text{Tr}(\rho_{\partial}(t) f(\Gamma[\cdot])), \quad (4.2)$$

where $d\mu$ is the natural measure induced on $\mathcal{X}$. Under the same short-memory and mesoscopic assumptions already used in Paper III, one obtains the generic continuity form

$$\frac{\partial p}{\partial t} + \nabla_A J^A = 0. \quad (4.3)$$

Equation (4.3) is not an optional phenomenological ansatz. It is the universal open-system shadow of the bulk-to-boundary projection once one chooses a finite observable dictionary.

The universalization step now consists of recognizing that the current J^A can always be decomposed into four conceptually distinct pieces,

$$J^A = J_{\text{grad}}^A + J_{\text{curl}}^A + J_{\text{noise}}^A + J_{\text{gauge}}^A, \quad (4.4)$$

with

$$J_{\text{grad}}^A = -\mu^{AB} \nabla_B \mathcal{F}_{\text{GFF}} p, \quad (4.5)$$

$$J_{\text{curl}}^A = \Omega^A p, \quad \nabla_A \Omega^A = 0, \quad (4.6)$$

$$J_{\text{noise}}^A = -D^{AB} \nabla_B p, \quad D^{AB} \succeq 0, \quad (4.7)$$

$$J_{\text{gauge}}^A = R^A_{\alpha} \lambda^{\alpha} p. \quad (4.8)$$

Here $\mathcal{F}_{\text{GFF}}$ is the effective geometric free-energy ledger that will later be developed systematically; μ^{AB} is the mobility tensor; Ω^A is the divergence-free circulation or curl flow; D^{AB} is the positive semidefinite diffusion tensor; R^A_{α} are the generators of dynamically maintained gauge directions; and $\lambda^{\alpha}(t, x)$ are the corresponding gauge velocities. The first three terms are the same gradient–curl–noise structure already derived in Paper III. The fourth is the new universal term that becomes explicit once gauge redundancy is recognized as a dynamical feature of open systems rather than merely a microscopic redundancy of field variables.

Equations (4.3)–(4.8) are the universal runtime kernel of the present paper. They say that every open system in nature can be described, at the correct level of coarse-graining, by an effective state manifold, a continuity law, a ledger functional, a solenoidal circulation, a fluctuation tensor, and a set of dynamical gauge directions. Paper IV-B later refines these ingredients into the SIP formalism, the GFF ledger, and the Qi-4 pillars; but mathematically the engine is already visible here.

Definition 4.1 (Universal open-system runtime packet). For the purposes of the present section, a *runtime packet* is the tuple

$$\mathfrak{S}_{\text{run}} = (\mathcal{X}, p, \mathcal{P}, \mathcal{O}, \mathcal{F}_{\text{GFF}}, \mu^{AB}, D^{AB}, \Omega^A, R^A_{\alpha}), \quad (4.9)$$

where $\mathcal{X}$ is the effective state manifold, p the reduced state density, $\mathcal{P}$ the set of ports through which exchange occurs, $\mathcal{O}$ the admissible observable dictionary, $\mathcal{F}_{\text{GFF}}$ the effective ledger functional, μ^{AB} and D^{AB} the dissipative and stochastic tensors, Ω^A the circulation field, and R^A_α the generators of gauge directions.

Remark 4.2. Equation (4.9) is intentionally provisional. Later sections will replace it by the more structured SIP formalism. The purpose of the present definition is only to make the universal extension of GDOS mathematically explicit before the more specialized vocabulary is introduced.

4.2 Why every scientific observation is constitutionally open

The central thesis of Paper IV-B is that openness is not an accidental imperfection of real experiments. It is an unavoidable structural feature of scientific observation itself.

To see this cleanly, consider a hypothetical closed law on a total state space $\mathcal{Z}$, written abstractly as

$$\dot{z} = \mathcal{V}(z), \quad z \in \mathcal{Z}. \quad (4.10)$$

A finite observer never has direct access to the full z . What is actually available is a map

$$\Pi : \mathcal{Z} \longrightarrow \mathcal{X}, \quad x = \Pi(z), \quad (4.11)$$

followed by an instrument map

$$\mathcal{M} : \mathcal{X} \longrightarrow \mathcal{Y}, \quad y = \mathcal{M}(x), \quad (4.12)$$

where $\mathcal{Y}$ is the finite record space: counts, images, voltages, trajectories, labels, symbols, reports, or numerical summaries. Unless Π is globally invertible and the instrument is noiseless and complete, the induced evolution on $\mathcal{X}$ and $\mathcal{Y}$ cannot be closed in the sense of Eq. (4.10). Memory kernels, noise, hidden variables, or control ports necessarily appear.

This fact may be stated as a theorem.

Theorem 4.3 (Universal openness of finite observation (T-Qi4-3.3)). *Let $\mathcal{Z}$ be a total state space with closed evolution $\dot{z} = \mathcal{V}(z)$. Let $\Pi : \mathcal{Z} \rightarrow \mathcal{X}$ be a non-invertible coarse-graining map and let $\mathcal{M} : \mathcal{X} \rightarrow \mathcal{Y}$ be a finite instrument map. Then the induced evolution on $\mathcal{X}$ is generically non-autonomous or non-Markovian unless additional hidden variables are appended; equivalently, the observer-accessible dynamics is an open dynamics.*

Proof. Because Π is non-invertible, there exist distinct total states $z_1 \neq z_2$ such that $\Pi(z_1) = \Pi(z_2) = x$. The future derivatives of the reduced coordinates are

$$\dot{x}_i = d\Pi_{z_i}[\mathcal{V}(z_i)], \quad i = 1, 2.$$

Generically these differ because $\mathcal{V}(z_1) \neq \mathcal{V}(z_2)$ and because the fibers of Π are not invariant under the total flow. Hence $\dot{x}$ cannot be written as a single-valued autonomous function of x alone. To close the reduced law one must either add memory, add hidden variables, or average over unresolved fibers.

All three operations are precisely the standard signatures of open-system dynamics. The additional instrument map $\mathcal{M}$ can only discard more information, never restore invertibility. Therefore the observer-accessible dynamics is open. ■

Theorem 4.3 is conceptually decisive. It implies that the quantum measurement problem is not the whole story but only the most historically visible instance of a much broader fact. Every field of science studies projected, instrument-conditioned, finite descriptions of larger systems. The difference between physics and biology is not that one is open and the other is closed. The difference is only the complexity of the projection and the richness of the resulting state manifold.

Table 5: How the universal open-system logic appears across scientific domains. Five domains run down the first column, from a particle experiment to a collective or social system. The second column gives the accessible coordinates in each case, the third the typical ports through which the system exchanges with what it cannot resolve, and the fourth why the induced dynamics is open rather than closed. The point of setting them side by side is that the fourth column reads the same in every row: the openness is not a feature of the subject matter but of the fact that a finite observer resolves only part of the state.

Domain	Accessible coordinates x^A	Typical ports	Why the induced dynamics is open
Particle experiment	Detector counts, event classes, low moments	Beam injection, environmental decoherence, readout electronics	Hidden bulk/environmental modes are traced over before a record is formed.
Chemical reactor	Concentrations, pH, temperature, flow rates	Inlet/outlet streams, heating, catalysts	Reservoirs and chemostats continuously exchange matter and energy with the reaction volume.
Cell or organism	Metabolite pools, membrane potentials, gene-expression states	Nutrient inflow, waste outflow, sensory coupling, action channels	The organism maintains itself only by continuous exchange across a boundary.
Neural circuit	Firing rates, phases, synaptic efficacies, bump positions	Sensory inputs, motor outputs, neuromodulation, blood flow	Hidden microstates and continual environmental coupling prevent closure on the recorded variables alone.
Collective or social system	Densities, prices, beliefs, network statistics	Communication channels, trade, migration, media, regulation	The coarse observables are projections of many unresolved agents and changing external constraints.

4.3 Ports, currents, and the minimal bookkeeping common to all systems

Once openness is accepted as universal, every system must be described not only by internal coordinates but also by its ports. A port is any channel across which the system exchanges quantities that matter for persistence: energy, matter, charge, entropy, information, stress, nutrients, toxins, prediction errors, or control signals. In the later SIP formalism this will be stated with greater precision, but the bookkeeping idea is already straightforward.

Let $X^I(t)$ denote any extensive quantity tracked at the effective level, with I ranging over the ledger sectors relevant to the problem. Then the generic balance law is

$$\frac{dX^I}{dt} = \sum_{a \in \mathcal{P}_{\text{in}}} J_a^I - \sum_{b \in \mathcal{P}_{\text{out}}} J_b^I + \Sigma_{\text{int}}^I, \quad (4.13)$$

where J_a^I are boundary currents and Σ_{int}^I denotes internal conversion between sectors. For strictly conserved quantities $\Sigma_{\text{int}}^I = 0$; for quantities like entropy or prediction error, Σ_{int}^I is generically nonzero.

Equation (4.13) is the common grammar behind systems that are often treated as if they were conceptually unrelated. In a biochemical network it describes nutrient and waste exchange. In a neuron it describes ionic, metabolic, and information-like currents. In ecology it describes immigration, predation, and resource uptake. In economics it describes capital, goods, and information exchange. What differs across domains is the semantic meaning of X^I , not the necessity of the balance law.

The special role of the GG-6/GDOS framework is that it does not let this bookkeeping remain an arbitrary phenomenological choice. Because the entire program began from a geometric gauge ledger in six dimensions, it insists that the bookkeeping for open systems must eventually be tied back to geometry and gauge structure. That is precisely why GFF, SIP, and then Qi-QUA become necessary in the final section order. But even before those later sections are written, the universalization of GDOS already tells us that every admissible open-system description must declare its ports and currents.

Proposition 4.4 (Runtime form of every admissible open-system model). *Within the GG-Theory/GG-6/GDOS program, every admissible effective model of an open system must specify:*

- (i) an effective state manifold $\mathcal{X}$,
- (ii) an observable dictionary $\mathcal{O}$,
- (iii) a port set $\mathcal{P}$,
- (iv) a continuity law of the form (4.3),
- (v) a current decomposition of the form (4.4),
- (vi) and a balance ledger of the form (4.13).

Proof. Items (i)–(iii) are forced by the universal openness theorem: a nontrivial observational description requires reduced coordinates, observables, and ports. Item (iv) follows from probability conservation on the effective manifold. Item (v) is the generic gradient–curl–noise decomposition augmented by the gauge directions made explicit in Paper IV-B. Item (vi) is the integrated form of the continuity equation applied to the extensive quantities of interest. Each item is therefore not an added decorative choice but a necessary consequence of treating scientific observation as open and lawful. ■

4.4 Why the new gauge term is unavoidable

At first glance the last term in Eq. (4.4) may appear optional. One can often write phenomenological open-system models with only gradient, curl, and noise. Why insist on the gauge term $J_{\text{gauge}}^A = R^A_\alpha \lambda^\alpha p$?

The answer is that Paper IV-B is not merely writing down useful phenomenology. It is asking for the most general form that remains faithful to the derivation chain from GG-6 and OGGEF. Paper III already showed that the observable boundary dynamics descends from a bulk with genuine gauge structure (Arisaka, 2026, A2-GG6; Arisaka, 2026, A5-GDOS). Once that is accepted, it is no longer consistent to treat open-system gauge structure as a negligible metaphor. The reduced system may not retain the full microscopic local gauge symmetry, but it does retain an induced equivalence relation under time-dependent boundary transformations. The new term in Eq. (4.8) is simply the mesoscopic imprint of that fact.

The mathematical meaning is simple. The generators R^A_α span directions in the effective state manifold that correspond to dynamically equivalent descriptions. Motion along these directions changes coordinates or internal phase conventions but not physical observables. If that redundancy were frozen once and for all, one would recover a conventional fixed-gauge formalism. But in an open system the environment continually changes the effective phase conventions, synchronization relations, and coordinate alignments. Therefore the gauge directions themselves become dynamical. This is the origin of DGI.

The practical importance is equally large. Without explicit gauge directions one cannot distinguish genuine adaptation from mere coordinate drift, genuine symmetry breaking from gauge fixing, or genuine attractors from artifacts of parameterization. In a theory that aims to cover systems from quantum devices to brains, that distinction is not optional.

Table 6: Why the universal GDOS kernel requires four current sectors rather than only three. The four — gradient, curl, noise and gauge — run down the first column, with their mathematical role, their physical meaning, and the reason each is indispensable in the columns that follow. The last column is the argument: dropping any one of the four does not merely coarsen the description, it breaks a conservation statement the boundary law needs, and the table names which one in each case.

Current sector	Mathematical role	Physical meaning	Why it is indispensable
Gradient	$-\mu^{AB}\nabla_B\mathcal{F}_{GFF}p$	Dissipative descent or relaxation	Captures irreversible tendency toward lower effective ledger cost.
Curl	$\Omega^A p, \nabla_A\Omega^A = 0$	Circulation, recurrent flow, Hamiltonian-like transport	Captures persistent loops, cycles, and driven recurrences without scalar descent.
Noise	$-D^{AB}\nabla_B p$	Fluctuation, uncertainty, coarse-graining residue	Captures stochasticity and hidden-variable dispersion inherent to finite observation.
Gauge	$R^A{}_\alpha\lambda^\alpha p$	Motion along dynamically equivalent descriptions	Captures the inherited gauge redundancy of open systems and prevents confusion between physical and coordinate change.

4.5 Section summary: what has now been universalized

The work of the present section may be summarized compactly.

First, the open-system law derived in Paper III has been detached from its original quantum-measurement presentation and rewritten in the generic language of reduced coordinates, probability currents, ports, and balance laws. This makes the GDOS kernel applicable in principle to any branch of science.

Second, the universal openness theorem has been proved in its basic form: whenever a finite observer studies a larger system through a non-invertible projection and a finite instrument, the accessible dynamics is open. The measurement problem in quantum mechanics is therefore one special case of a universal observation problem.

Third, the mesoscopic current decomposition has been sharpened from the three-sector form of Paper III to the four-sector form required for Paper IV-B. The new gauge sector makes it possible to speak rigorously about dynamically maintained equivalence classes of open-system descriptions. That step is the doorway to DGI.

Fourth, the universal runtime packet (4.9) has been declared. It is intentionally provisional, because the later SIP formalism will refine it. But it is sufficient to support the technical derivation of the master pipeline.

The next section therefore performs the first major theorem-level construction of Paper IV-B: it derives the runtime pipeline

$$\text{lawful open-system evolution} \implies \text{DGI} \implies \text{GSSB} \implies \text{GA}.$$

That pipeline is the minimal grammar by which open systems search, select, and stabilize.

Bridge to Part II. With the entrance and recap of Part I in hand — the constitutional inputs from Arisaka (2026, A1-Constitution), the 6D bulk and BI-6 closure from Arisaka (2026, A2-GG6) and Arisaka (2026, A3-BI6), the GDOS boundary law from Arisaka (2026, A5-GDOS), and the four geometric pillars TGG-CCI-GFF-GRIP from Arisaka (2026, A6-GRIP) — we have everything we need to begin the Qi-runtime construction. Part II imports the runtime constitution: GRIP as the cascade structure, GFF as the master functional, and SIP as the canonical unit of analysis. These three are not derived in the present paper; they are the standing inputs from GG-A6-GRIP on which the four Qi-pillars of Parts III–VI will be built.

5 The cohort inheritance and the roadmap of the fifteen Parts

Before the imported runtime constitution begins (Part II), it is useful to lay out two reader-orientation maps on the same opening: a panorama of the twenty-two-paper cohort in which this paper stands (§5.1, Figure 2), and a panorama of the fifteen Parts of the present paper (§5.2, Figure 3). A reader of any later Part should be able to flip back to this Section and orient themselves in both directions: which cohort companion carries which imported result, and where the present argument stands in the fifteen-Part arc.

5.1 The cohort inheritance

The present paper is one of twenty-two. GG-Theory is written not as a single monograph but as a cohort of twenty-two preprints — eleven A-papers, six P-papers, three C-papers and two white papers — each carrying one rigorously stated piece of a single descent that runs from four axioms and seven postulates to the empirical structure of physics. The partition into twenty-two papers reflects the partition of subject matter, not a ranking of importance: the framework is one program, and each paper is a named carrier of named content. Figure 2 is the panorama; a reader meeting the program for the first time can read its architecture off the figure top to bottom before returning to the information-theoretic runtime that is the business of this paper. The cohort falls into five strata plus three framed branches.

The constitutional root and the GG6 Trio. At the root stands GG-A1-Constitution (Arisaka, 2026, A1-Constitution), the GG-Constitution: the four axioms (AX-1 LC, AX-2 GGEP, AX-3 OGGE, AX-4 MSMF) and seven postulates against which every companion’s theorems are audited. Forced from it is the GG6 Trio that builds the bulk: GG-A2-GG6 (Arisaka, 2026, A2-GG6) constructs the six-dimensional Geometry–Gauge bulk; GG-A3-BI6 (Arisaka, 2026, A3-BI6) proves the modified

Bianchi identity BI-6 anomaly-free; and GG-A4-Euclid ([Arisaka, 2026, A4-Euclid](#)) traces the lineage from Euclid to the 6D bulk and reconstructs the universal constants.

The boundary law and its two pillars — where this paper sits. AX-3 (OGGEP) forces the open-system boundary law GG-A5-GDOS ([Arisaka, 2026, A5-GDOS](#)), the Geometric Dynamics of Open Systems, which projects bulk physics onto the four-dimensional boundary. Two sibling pillar papers organize the boundary content, and *the present paper is one of them*: GG-A7-Qi4 ([Arisaka, 2026, A7-Qi4](#)) — this paper — supplies the information-theoretic Qi-4 pillars (Qi-QUA, Qi-EP, Qi-CON, Qi-HAL), the runtime grammar carried by the memory-bearing $U(1)_{Qi}$ sector of the boundary; its geometric sibling GG-A6-GRIP ([Arisaka, 2026, A6-GRIP](#)) supplies the Gi-4 pillars (TGG, CCI, GFF, GRIP), and with the two pillar papers in hand GG-A5-GDOS closes the GDOS Trio.

The Consequence Trio and the capstone. Three papers read the empirical world off the boundary law: GG-A8-QM ([Arisaka, 2026, A8-QM](#)) recasts four-dimensional quantum mechanics, GG-A9-SM ([Arisaka, 2026, A9-SM](#)) derives the twenty-eight free parameters of the Standard Model from the OGN-5 ledger and the imported Gi-4/Qi-4 runtime, and GG-A10-Cosmos ([Arisaka, 2026, A10-Cosmos](#)) derives the Λ CDM cosmological inventory; GG-A11-Origin ([Arisaka, 2026, A11-Origin](#)) is the capstone that closes the A-series and generalizes the ground-attractor statement across the layers of the program-wide cascade. The per-stratum statement on which the present paper’s GGA terminus rests is upstream, in GG-A6-GRIP ([Arisaka, 2026, A6-GRIP](#)).

The two phenomenology branches. Two framed branches carry the sharp predictions. The Predictions Branch (P-series) collects the KK-resonance and hierarchy paper GG-P1-KK ([Arisaka, 2026, P1-KK](#)), the three-generations paper GG-P2-N3 ([Arisaka, 2026, P2-N3](#)), the electron paper GG-P3-Electron ([Arisaka, 2026, P3-Electron](#)), and the proton paper GG-P6-Proton ([Arisaka, 2026, P6-Proton](#)), whose baryon number is a winding of the same compact circle this paper tracks as the carrier of information. The Cosmology & Astrophysics Branch (C-series) collects the Kaxion dark-matter paper GG-C1-Kaxion ([Arisaka, 2026, C1-Kaxion](#)) — the high-energy carrier of the same $U(1)_{Qi}$ that the present paper tracks at chemical and biological scales — the Hubble-tension paper GG-C2-Hubble ([Arisaka, 2026, C2-Hubble](#)), and the black-hole-entropy paper GG-C3-BH ([Arisaka, 2026, C3-BH](#)).

The Foundations Branch. The third framed branch is the Foundations Branch: the two member papers whose subject is the boundary’s complex structure and its mirrors. GG-P4-Real ([Arisaka, 2026, P4-Real](#)) stands on the map’s left flank as the origin-of- i member, deriving the single imaginary unit of boundary physics from two real structures of the GG-A2-GG6 arena — the compact phase circle and the causal clock. For the present paper its significance is direct: GG-A7-Qi4 carries the four Qi-4 pillars (Qi-QUA, Qi-EP, Qi-CON, Qi-HAL) — the quantum-information runtime that runs on the memory-bearing $U(1)_{Qi}$ systems of the boundary — and GG-P4-Real is the origin of the boundary’s complex structure whose phase the $U(1)_{Qi}$ memory carrier stores. GG-P5-CPT ([Arisaka, 2026, P5-CPT](#)) carries the three mirrors of matter — C, P, and T read as orientation reversals of the three real bulk factors — and its law-level arrow of time is kin of the one-way information

cascade that the runtime built here drives in a single direction. Both carry P-series designators; their place at the spine’s flank, not the phenomenology rows, marks their foundational role.

The White-Paper Branch. A fourth group stands on the map’s right flank and reads the spine rather than feeding it. GG-W1-Symmetry ([Arisaka, 2026](#), [W1-Symmetry](#)) walks the century from Noether’s theorem to the differential-cohomological statement of anomaly cancellation, and GG-W2-Holonomy ([Arisaka, 2026](#), [W2-Holonomy](#)) dissects what that arc arrives at — protection by topology rather than by symmetry. Both bear on the present paper at the same point: the memory this paper’s fourth pillar formalizes is holonomy-class data on a compact circle, and a quantity protected by a winding is protected in the way those two papers describe, not in the way a conserved charge is.

The twenty-two-paper map, identical in every paper of the cohort, is Figure 2, which closes this subsection.

5.2 Roadmap of the fifteen Parts

The rest of the paper is organized into fourteen further Parts, and the walk through them is plotted in Figure 3. The present Part — the entrance and bridge — is nearly done: the strategic entrance, the recap of the constitutional chain imported from the companions, the cohort the paper stands in, and the universal extension of GDOS to any open system in nature. Everything after this page is the runtime itself, and the itinerary below says where each piece lives.

The entrance conditions come first. Part II imports the runtime constitution from GG-A6-GRIP in the strict order demanded by logical dependence: first the master pipeline $\text{GRIP} = \text{DGI} \rightarrow \text{GSSB} \rightarrow \text{GA}$, then the scalar ledger GFF, and only then the canonical system-instance packet SIP — with the cascade-closure import flagged conditional on the companion’s cascade conditioning (M11, in spectral-flow form). Nothing informational has been claimed yet; Part II simply sets the geometric stage on which information will be made physical.

The six core Parts are the apex, and they are highlighted in the roadmap because together they deliver the paper’s central claim. The four Qi-4 pillars are established there as principal theorems in their own dependency order: Qi-QUA locks the informational current to the energetic, electric, and baryonic currents through the Arisaka equation and its single conversion constant Q_B ; Qi-EP gives the three extremal traditions their histories at review depth, raises the variational principle that selects the runtime’s trajectories, derives the three gauges in which the selection is read, and derives also the boundary past which those readings must fail; Qi-CON derives the conformal, Lévy, fractal, and $1/f$ diagnostics by which the runtime is recognized in data; and Qi-HAL builds the memory-bearing attractor architecture on the $U(1)_{Q_i}$ sector. Part VII then adds the persistence and completeness layer (SeOS, QCI) — what holds itself together, and how much of the runtime a given system actually runs — and Part VIII closes the block with the statement that the ground attractor of the geometric companion carries all four pillar certificates, proved in that direction with the converse recorded as open. Every earlier Part prepares that construction, and every later Part spends it.

Part IX builds the internality layer on top of the pillars: the Outer Hologram, the Inner Hologram, and their closed-loop coupling $\text{OH}+\text{IH}=\text{EIC}$ — the internal-measurement runtime by which an open system comes to carry a model of its world — together with the EIC universality theorem and the graded continuum up to consciousness-grade runtime.

Three Parts then widen and weigh the case. Part X treats emergence and scope: the Japanese five-story pagoda hierarchy and the cross-domain applications of the runtime to chemistry, life, brains, and society, with the explicit universality claim. Part XI is the assessment: significance and uniqueness, the cross-paper coherence map, the triaged open problems, and the joint falsifiability program. Part XII draws the conclusions: the synthesis of the Qi-4 runtime and the bargain with the laboratory.

The back matter is the workshop. Part XIII collects the technical derivations, the cross-domain worked examples, and the abbreviation, symbol, and key-equation references; Part XIV answers the strongest objections a hostile particle theorist, systems biologist, origin-of-life researcher, neuroscientist, and methodologist would raise, with pointers into the paper; and Part XV closes with the annotated references, each carrying a Relevance note. Read in order, the fifteen Parts




Figure 2: The twenty-two-paper cohort map of GG-Theory. The constitutional root **GG-A1-Constitution** (violet, top) supplies the four axioms and seven postulates from which the rest of the cohort descends: the GG6 Trio (**GG-A2-GG6**, **GG-A3-BI6**, **GG-A4-Euclid**, red) realizes the bulk geometry; **GG-A5-GDOS** (yellow) the boundary law; the GDOS Trio pillars (**GG-A6-GRIP**, **GG-A7-Qi4**, green) the geometric and quantum-information runtime; the 4D-Physics Trio (**GG-A8-QM**, **GG-A9-SM**, **GG-A10-Cosmos**, blue) the textbook readouts; and the closing capstone **GG-A11-Origin** (gold) the quarks-to-consciousness identity cascade. Two framed branches carry the phenomenology — the **Predictions Branch** (cyan), whose four papers hand their results along the row by the internal arrows (**GG-P2-N3**’s $N_g=3$ into **GG-P3-Electron**, and that paper’s ground-attractor method, *GGA*, on into **GG-P6-Proton**), and the **Cosmology & Astrophysics Branch** (teal). The spine is flanked by two columns of two, which do opposite jobs. On the left the **Foundations Branch** (violet) feeds it: **GG-P4-Real**, fed by **GG-A2-GG6**’s Theorem 0 arena, manufactures the i that **GG-A8-QM** consumes, and **GG-P5-CPT** is the **GG-A9-SM**–**GG-A10-Cosmos** interface, fed by the ledger phases and feeding the η_B chain across the width of the map, with the shared- i edge descending solid between them. On the right the **White-Paper Branch** (brown) reads it: **GG-W1-Symmetry** walks the century from Noether to Hopkins–Singer, and **GG-W2-Holonomy** dissects the structure that arc arrives at — protection by topology rather than by symmetry — across sixteen open problems. The white papers derive nothing new, carrying every result at its parent’s declared grade; the single dashed feed drawn into each names its heaviest source, not its only one. *Reading conventions:* solid arrows mark axiom-forced descent and carry an edge label naming the driving axiom or postulate where most illustrative; dashed arrows mark the deployment of derived results, or historical links inherited without being constitutionally forced.

compose into a single self-contained construction that any reader can audit step by step; read by need, each Part names its imports and can stand a visit on its own.




Figure 3: Roadmap of the fifteen Parts of GG-A7-Qi4. The argument develops in five phases. *Foundations and the imported runtime* (Parts I–II): the entrance and bridge, then the runtime constitution GRIP, GFF, and SIP imported from GG-A6-GRIP, whose closure is conditional on that companion’s remaining cascade conditioning — the monodromy statement M11 in spectral-flow form, with the budget question M10 resolved externally and the scope refinements closed at class level by the three-generations companion. *The Qi-4 core* (Parts III–VIII), the six banded Parts: the four informational pillars established as principal theorems in their dependency order ($Qi-QUA \Rightarrow Qi-EP \Rightarrow Qi-CON \Rightarrow Qi-HAL$), then the persistence and completeness layer (SeOS, QCI), then the statement that the ground attractor of the geometric companion carries all four pillar certificates — proved in that direction, with the converse recorded as open. They are drawn as one block because they are one argument told in six instalments, and the block is highlighted because it is the paper’s central claim. *Internal measurement* (Part IX) builds the outer and inner holograms and the loop that couples them on top of the pillars. *Emergence and assessment* (Parts X–XI): the five-story pagoda and the cross-domain recipe, then significance, open problems, and falsifiability. The *conclusions* (Part XII) close the argument, and the three back-matter Parts collect the technical appendices (Part XIII), the common objections and responses (Part XIV), and the annotated references (Part XV).

6 Why none of this is a special case

The result this Part rests on is easy to state and easy to underestimate. Openness is not a property that complicated things have and simple things lack. It is the price of being finite. Any observer who cannot hold the whole state of a system — and no observer can — is handed a description that will not close on itself, and must carry memory, or hidden variables, or averaging to make up the difference. That is as true of a particle detector as of a rainforest.

It follows that the boundaries between the sciences are not boundaries in nature. They are differences in how much survives a projection. Particle physics studies a projection so severe that only a handful of numbers come through, which is why its laws can be written on a page. Ecology studies one through which so much comes through that no page would hold it. Neither is more or less open than the other, and the habit of treating the first as fundamental and the second as merely complicated has no support in the structure. It is a statement about how much of the world each field agreed to look at.

There is a consequence for the word “fundamental” that is worth drawing out, because that word does a great deal of unexamined work. On the usual picture the fundamental science is the one that studies the smallest things, and every other science is an account of what happens when many of them are put together. On the picture this Part argues for, “fundamental” would have to mean the level at which a description closes on itself — and that level is nowhere. No science studies a closed system, because there are none to study. What separates the disciplines is not distance from a foundation but the severity of the projection each has agreed to work behind. That is not a demotion of the physics of small things. It is a promotion of everything else to the same standing, which is a different claim and a more interesting one.

There is a second thing worth taking from this Part, and it is about observation rather than about systems. It is tempting to think that the act of measuring is what makes a system open — that the world is whole until someone looks at it. The framework says the opposite. The projection is running everywhere, at every instant, whether or not anyone is studying anything. An observer does not open the world. An observer inherits a world that was already open, and inherits it in the particular, partial, instrument-shaped way that being one small thing inside it requires.

That is a humbling thought and also a liberating one. Humbling, because it means the view from anywhere is a view from somewhere, and no apparatus will ever escape into the unprojected. Liberating, because it means the partiality is lawful. It has a form, the form can be written down, and what follows in this paper is an attempt to write it.

It is worth being clear that none of this is resignation about the limits of knowledge. The claim is not that the world is unknowable and one ought to be modest about it; that sentiment is cheap and commits nobody to anything. It is sharper. The incompleteness of every description has a definite form, the form can be computed, and what a projection discards can be accounted for rather than regretted. A limit one can write down is not the same kind of object as a limit one can only gesture toward: the first can be carried into a calculation and argued about by people who disagree. Most of what this paper does from here on is that — taking a thing usually said in the tone of a caution and insisting on saying it as a quantity.

Part II

Runtime constitution imported from GG-A6-GRIP

GRIP, GFF, and SIP as the entrance conditions for the Qi-runtime

This Part is a customs declaration. It carries nothing new, and it is here because what follows would be unreadable without it.

Three things have to be on the table before a single sentence about information can be written down. The first is an engine: a system holds all its options open, then loses stability and commits to one of them, then settles somewhere it can stay. Those three moves are the whole mechanism by which anything in nature comes to be one definite thing rather than a smear of possibilities, and they are established in the companion paper on the geometric side, not here. The second is a budget. Exploring, committing and holding still are three activities that have to be paid for out of one account, and the companion supplies that single ledger together with the proof that it is single — there is not one free energy for the chemistry and another for the information. The third is a standard way of writing down one concrete system, so that a protocell, an enzyme, a navigating insect and an institution can be set side by side without pretending they are alike.

Two points in the import deserve attention even from a reader who plans to move quickly. The first is that the destination of the engine is not merely a stable state but the deepest one available on its stratum, and that distinction is why matter is permanent rather than merely long-lived. The second is that the engine runs in two regimes, and the difference between them is a single question: does the landscape itself move? Descent on a landscape that stays put finishes a journey. The transient that reorganizes the landscape starts one. Cosmology supplies the second kind from outside, as unrepeatable events; and the discovery that a system with a memory can produce them from inside, on purpose, is what this paper adds later.

One theorem in this Part is proved here rather than imported, and it is worth knowing which. It says that the standard description of a system contains neither too little nor too much: everything an experiment could return is fixed by it, and two descriptions that predict the same things differ only by a relabeling. The fifth and last field of that description is the one to notice — it holds the criteria under which the system's own model would be declared wrong. Falsifiability is built into what it means to write a system down here, rather than added at the end.

Nothing informational has been claimed by the close of this Part. It sets the stage; the next Part puts something on it.

7 The imported runtime theorem: GRIP

Part I established the constitutional bridge from Papers I–IV-A to a universal theory of open systems. Part II now begins the constructive runtime sequence with the dynamical grammar imported from GG-A6-GRIP: lawful exploration (DGI), branch selection (GSSB), and attractor stabilization (GA). Only once this three-stage engine is explicit does it become possible to ask which scalar ledger it descends (GFF) and which concrete instance carries it (SIP). The master pipeline takes the boxed form

$$\boxed{\text{DGI : organize equivalences}} \implies \boxed{\text{GSSB : select branch}} \implies \boxed{\text{GA : stabilize outcome}} \quad (7.1)$$

each ingredient strictly needed: DGI without GSSB gives lawful exploration without commitment; GSSB without DGI risks mistaking coordinate choices for physical commitments; DGI+GSSB without GA gives no persistence criterion.

DGI at the operator level. Let the boundary-visible density obey $d\rho_\partial/dt = \mathfrak{L}_t[\rho_\partial]$ and let U_t be a time-dependent unitary family induced from an admissible bulk gauge transformation under OGGEF. Standard manipulation gives the generator-transformation law

$$\mathfrak{L}_t^{(U)} = \mathcal{U}_t \circ \mathfrak{L}_t \circ \mathcal{U}_t^{-1} - i[K_t, \cdot], \quad \mathcal{U}_t[\rho] = U_t \rho U_t^\dagger, \quad K_t := i \dot{U}_t U_t^\dagger, \quad (7.2)$$

which is the exact operator-level statement of DGI: the reduced law is defined by an equivalence class of generators related by induced gauge transformations (Arisaka, 2026, A5-GDOS, Breuer and Petruccione, 2002).

GSSB on the gauge-reduced manifold. On the physical quotient $\mathcal{X}_{\text{phys}}$ a symmetric invariant branch can lose stability and lower-symmetry branches appear. The local Landau picture $V_\eta(q) = \frac{1}{2}a(\eta)q^2 + \frac{1}{4}b(\eta)q^4$ with bifurcation at $a(\eta) = 0$ is the canonical model; the open-system version replaces deterministic minima with stationary distributions that concentrate on broken branches when $a(\eta) < 0$ (Touchette, 2009, van Kampen, 2007).

GA as the selected persistent invariant.

Definition 7.1 (Geometric Attractor). A *Geometric Attractor* $\mathcal{A}_{\text{GA}} \subseteq \mathcal{X}_{\text{phys}}$ is an invariant set or invariant measure satisfying: (i) invariance under the reduced GDOS semigroup; (ii) asymptotic relevance (a nontrivial basin converges to it); (iii) selection by GSSB (it sits on a symmetry-reduced branch from a larger admissible family); (iv) gauge admissibility (defined on the physical quotient); (v) ledger admissibility (compatible with the GFF descent or steady-balance constraints).

Theorem 7.2 (Existence of invariant asymptotic sets under confining GFF landscapes (T-Qi4-6.2)). *Consider the reduced flow*

$$\dot{y}^i = -\tilde{\mu}^{ij} \partial_j \tilde{\mathcal{F}}_{\text{GFF}} + \tilde{\Omega}^i, \quad (7.3)$$

on $\mathcal{X}_{\text{phys}}$, with $\tilde{\mu}^{ij}$ positive semidefinite, $\tilde{\Omega}^i \partial_i \tilde{\mathcal{F}}_{\text{GFF}} = 0$, and $\tilde{\mathcal{F}}_{\text{GFF}}$ proper and bounded below. Then every trajectory is contained in a compact sublevel set of $\tilde{\mathcal{F}}_{\text{GFF}}$, and its ω -limit set lies in the largest invariant subset of $\mathcal{C} = \{y \in \mathcal{X}_{\text{phys}} \mid \partial_i \tilde{\mathcal{F}}_{\text{GFF}} \tilde{\mu}^{ij} \partial_j \tilde{\mathcal{F}}_{\text{GFF}} = 0\}$.

Proof. See GG-A6-GRIP Theorem T4 (GRIP principal theorem). ■

Theorem 7.3 (Master runtime pipeline (T-Qi4-6.3)). Assume (i) the GG-6 bulk and OGGEF reduction of Papers II and III; (ii) the universal open-system runtime kernel (4.3)–(4.8); (iii) a proper gauge-invariant effective ledger $\mathcal{F}_{\text{GFF}}$; (iv) a symmetry group $\mathcal{S}$ acting on $\mathcal{X}_{\text{phys}}$; and (v) a control regime in which the symmetric branch loses stability. Then the open-system dynamics factors as

$$\mathcal{X} \xrightarrow{\text{quotient}} \mathcal{X}_{\text{phys}} \xrightarrow{\text{instability}} \text{symmetry-broken branch} \xrightarrow{\text{asymptotic selection}} \mathcal{A}_{\text{GA}}. \quad (7.4)$$

Equivalently: induced gauge covariance yields DGI; loss of stability on the quotient yields GSSB; asymptotic concentration on the selected symmetry-broken invariant set yields a GA.

Proof. See GG-A6-GRIP Theorem T4 (GRIP principal theorem, dual-layer compatibility proposition $\mathfrak{C}_{\text{GRIP}}$). ■

Example: the Higgs vacuum as the particle-physics limiting case. The standard Higgs mechanism is the limiting case in which the open-system language collapses to familiar field-theory symmetry breaking (Higgs, 1964). With $V(H) = -\mu_H^2 |H|^2 + \lambda_H |H|^4$ and $\mu_H^2, \lambda_H > 0$, the gauge-equivalence class of vacua selects one orbit representative modulo the residual $U(1)_{\text{EM}}$: an elementary GSSB event whose GA is the symmetry-broken electroweak vacuum manifold. Paper IV-B does not replace this mechanism; it recognizes it as one corner of a much larger runtime grammar. Further worked examples (laser order parameter, ATP synthase, neural decision) appear in Appendix D and in the per-pillar sections.

Table 7: Operational audit of the master runtime pipeline. The three phases of the imported cascade — DGI, GSSB and the attractor — run down the first column. The second gives the mathematical test that identifies each phase, the third what changes as the system passes through it, and the fourth what must stay invariant across it. The fourth column is the one to check against an experiment: it lists the quantities whose violation would show that a system is not running this pipeline at all.

Phase	Mathematical test	What changes	What must stay invariant
DGI	Quotient by G_{dyn}	Coordinates, phases, internal alignment conventions	Gauge-invariant observables and physical predictions
GSSB	Instability of symmetric branch; isotropy reduction $H \subsetneq \mathcal{S}$	Branch, phase, architecture, committed mode	Underlying law and admissibility filters
GA	Asymptotic concentration on selected invariant set	Persistent pattern, cycle, measure, or policy	Compatibility with gauge reduction and ledger constraints

The terminus refined: from GA to GGA. Two completions of this imported grammar, both established at runtime level in GG-A6-GRIP (Arisaka, 2026, A6-GRIP), are adopted here and used throughout. The first refines the terminus. A GFF landscape generally admits *many* attractors satisfying Definition 7.1; the cascade output is not “some stable basin” but the unique deepest one.

Definition 7.4 (GGA: Ground Geometric Attractor; after GG-A6-GRIP). Among the Geometric Attractors of a stratum, the *Ground Geometric Attractor* is the global minimum of $\mathcal{F}_{\text{GFF}}$ on that stratum. The definition is GG-A6-GRIP’s and is imported unchanged (Arisaka, 2026, A6-GRIP); the fully resolved pipeline (7.4) therefore terminates not merely on a GA but on the GGA. Existence and uniqueness are GG-A6-GRIP’s per-stratum ground-state corollary, which holds on the class of strata that paper calls admissible and which carries that paper’s grade. No Qi-4 filter enters the definition: what the four pillars add is a *characterization* of the same object, and that is the business of the dedicated Part before Part IX (Part VIII).

The two regimes: GRIP-In and GRIP-Out. The second completion classifies the dynamics. The runtime operates in exactly two regimes, split by one question: *does the landscape change?* GRIP-In (GG-A6-GRIP) is all dynamics on a *fixed* landscape — the gradient flow (7.3) within a basin, instanton steps across the GSSB saddles between basins — halting uniquely at the stratum’s GGA. GRIP-Out (GG-A6-GRIP) is the complementary regime: a transient out-of-equilibrium event that *reorganizes the landscape itself* (the control parameters of $\mathcal{F}_{\text{GFF}}$ change) and supplies the initial datum of the next descent. In one sentence: GRIP-Out opens the landscape; GRIP-In finishes the journey. The entire master pipeline (7.1) is the GRIP-In regime; where the landscape itself came from is a GRIP-Out question, answered cosmologically in Arisaka (2026, A10-Cosmos) and layer by layer in GG-A11-Origin. What the present paper adds — and no other paper of the series can add — is the discovery that memory-bearing open systems *internalize* the second regime: the Qi-HAL GRIP-loop (§52) is a self-generated, recoverable GRIP-Out under the system’s own control.

Bridge to GFF and SIP. The remaining order is now fixed. Section 8 derives the scalar ledger GFF: the pipeline needs a scalar budget because exploration, commitment, and stabilization would otherwise have no common quantity to descend, replenish, or audit. Section 9 introduces SIP: once the ledger is fixed, the theory needs a concrete carrier with ports, generator, observables, and kill criteria. Parts III–VI then develop the four Qi pillars in their dependency order.

8 GFF: the geometric free-energy functional $\text{GFF} = \text{GEP} + \text{CEP}$

With GRIP in place the runtime needs a scalar ledger: a single quantity for the gradient sector to descend, the maintenance sector to replenish, and the curl sector to circulate around. MSMF forbids admitting several unrelated free-energy-like objects for the same packet, so the ledger is unique up to additive constants.

Definition 8.1 (The GFF functional). Given a packet with state manifold $(\mathcal{Z}, g_{AB})$, probability density p , and port configuration $\mathcal{P}$, the GFF ledger is

$$\mathcal{F}_{\text{GFF}}[p; \mathcal{P}] = \underbrace{\int_{\mathcal{Z}} p(x) \Phi_{\text{GEP}}(x; \mathcal{P}) d\mu_g(x)}_{\mathcal{F}_{\text{GEP}}} + \underbrace{\Theta_{\text{eff}} \int_{\mathcal{Z}} p(x) \ln \frac{p(x)}{p_{\text{ref}}(x; \mathcal{P})} d\mu_g(x)}_{\mathcal{F}_{\text{CEP}}}, \quad (8.1)$$

where $d\mu_g = \sqrt{\det g} d^n x$, Φ_{GEP} is the energy–geometry–gauge potential induced by the projected bulk structure, p_{ref} is the reference density of the induced geometric measure or prior boundary condition, and Θ_{eff} is the effective information-temperature scale.

The first term collects projected curvature costs, gauge potentials, holonomy contributions, binding energies, and field couplings (the GEP sector). The second is a relative-entropy/curvature-of-measure term capturing the open-system price of incomplete access (the CEP sector); recovering the standard sign convention gives $-\Theta_{\text{eff}} S$ plus an additive constant.

The full reduced current decomposes as

$$J^A = J_{\text{grad}}^A + J_{\text{curl}}^A + J_{\text{noise}}^A + J_{\text{gauge}}^A, \quad (8.2)$$

with $J_{\text{grad}}^A = -p M^{AB} \nabla_B (\delta \mathcal{F}_{\text{GFF}} / \delta p)$, $J_{\text{curl}}^A = \Omega^A p$, and $J_{\text{gauge}}^A = R^A_{\alpha} \lambda^{\alpha} p$. Substituting continuity into $d\mathcal{F}_{\text{GFF}}/dt$ and integrating by parts gives the balance identity

$$\frac{d}{dt} \mathcal{F}_{\text{GFF}} = -\mathcal{D}_{\text{GFF}} + \mathcal{P}_{\text{port}} + \mathcal{C}_{\text{curl}} + \mathcal{G}_{\text{gauge}}, \quad (8.3)$$

with dissipation $\mathcal{D}_{\text{GFF}} \geq 0$ quadratic in $\nabla(\delta \mathcal{F}_{\text{GFF}} / \delta p)$, port power $\mathcal{P}_{\text{port}}$ collecting explicit driving and boundary flux, and curl/gauge cross-terms

$$\mathcal{G}_{\text{gauge}} := \int_{\mathcal{Z}} \nabla_A \left(\frac{\delta \mathcal{F}_{\text{GFF}}}{\delta p} \right) R^A_{\alpha} \lambda^{\alpha} p d\mu_g, \quad (8.4)$$

together with $\mathcal{C}_{\text{curl}}$ of identical form with $\Omega^A p$. Both vanish whenever the curl current is tangent to $\mathcal{F}_{\text{GFF}}$ level sets and the ledger is gauge-invariant.

Theorem 8.2 (GFF monotonicity in the lawful sector (T-Qi4-7.2)). *Suppose $\mathcal{F}_{\text{GFF}}$ is gauge-invariant and $\Omega^A \nabla_A (\delta \mathcal{F}_{\text{GFF}} / \delta p) = R^A_{\alpha} \nabla_A (\delta \mathcal{F}_{\text{GFF}} / \delta p) = 0$. Then $d\mathcal{F}_{\text{GFF}}/dt = -\mathcal{D}_{\text{GFF}} + \mathcal{P}_{\text{port}}$; in the undriven closed limit the ledger is nonincreasing, and at a nonequilibrium steady state $\langle \mathcal{P}_{\text{port}} \rangle_{\tau} = \langle \mathcal{D}_{\text{GFF}} \rangle_{\tau}$.*

Proof. See GG-A6-GRIP Theorem T3 (GFF principal theorem, monotonicity in the lawful sector with compatibility proposition $\mathfrak{C}_{\text{GFF}}$). ■

The theorem disentangles three often-confused ideas. The system need not minimize GFF instantaneously; a driven open system may sustain nonzero circulation and throughput. Nevertheless, its gradient motion is organized by one ledger, and steady maintenance requires exactly enough port power to balance the irreversible descent and repair costs. This is the precise sense in which GFF generalizes both free-energy descent and nonequilibrium maintenance.

Remark 8.3 (Which Lyapunov function, at which level). The word “Lyapunov” is used for the runtime’s descent at four levels in this program, and the theorem above sits at a definite one. That *some* function decreases along the descent is Conley’s theorem (Conley, 1978): every semiflow on a compact state space has a complete Lyapunov function, with no hypothesis about the generator. That the function is the relative entropy to the layer’s stationary state is Lindblad’s and Spohn’s theorem for CPTP semigroups (Lindblad, 1975, Spohn, 1978), which needs a stationary state and nothing about detailed balance — a driven layer with net currents still descends, and what it descends is the relative entropy to its own steady state. That the function is $\mathcal{F}_{\text{GFF}}$ read as a free energy of the physical Hamiltonian, the second term of (8.1) with p_{ref} a Gibbs state, needs detailed balance with respect to that state (Kossakowski et al., 1977), which is the hypothesis of GG-A6-GRIP’s Theorem T3.2 and which a thermal, weakly coupled bath supplies. And that the function on an admissible stratum has one floor is Frankel’s and Atiyah’s (Atiyah, 1982, Frankel, 1959). Theorem 8.2 is at the second level: it says which terms the balance carries, and its undriven limit is the relative-entropy descent; the identification of Θ_{eff} with a temperature and of p_{ref} with a Gibbs state is the third level and is imported at GG-A6-GRIP’s grade. The Lyapunov functionals this paper builds for its own objects — the EIC functional, the SeOS viability bound, the composite pagoda functional — are model-level constructions and carry their own hypotheses where they are stated.

Once Qi-QUA is derived (Part III), the event-level vector ledger collapses to a one-parameter locked ray and scalarizes consistently with GFF; this is recorded in the event-budget appendix.

9 SIP: System Instance Packet as the Canonical Unit of GDOS

The scalar ledger of Section 8 fixes what is spent and replenished, but not on which carrier. Section 9 closes Part II by declaring the canonical packet: the minimal complete object on which the Qi pillars subsequently act.

Definition 9.1 (Canonical five-field SIP). For a concrete open-system instance in the GG-Theory/GG-6/GDOS program, the *System Instance Packet* is the ordered five-field object

$$\text{SIP} = \left[(\mathcal{Z}, g_{AB}), \mathcal{P}, (\mathcal{L}_{\text{GDOS}}, \eta), (\mathcal{F}_{\text{GFF}}, \mathcal{J}), (\mathcal{O}, \mathfrak{K}) \right], \quad (9.1)$$

where (i) $(\mathcal{Z}, g_{AB})$ is the effective state manifold and its metric; (ii) $\mathcal{P}$ is the full set of ports and imposed boundary couplings; (iii) $(\mathcal{L}_{\text{GDOS}}, \eta)$ is the lawful open-system generator together with its noise closure; (iv) $(\mathcal{F}_{\text{GFF}}, \mathcal{J})$ is the ledger bundle (scalar GFF plus the locked current vector $\mathcal{J} = (J_e, J_B, J_{Qi}, J_E)$); and (v) $(\mathcal{O}, \mathfrak{K})$ is the observable dictionary plus the explicit kill criteria required for falsifiability.

The packet plays, for open-system science, the role that metric-plus-matter plays in general relativity or that gauge-group-plus-matter-representation plays in Yang–Mills theory. The equivalent unfolded form

$$\widehat{\text{SIP}} = (\mathcal{Z}, g_{AB}; \mathcal{P}; M^{AB}, \Omega^A, D^{AB}, R^A_{\alpha} \lambda^{\alpha}; \mathcal{F}_{\text{GFF}}, \mathcal{J}; \mathcal{O}, \mathfrak{K}), \quad (9.2)$$

exposes drift, curl, diffusion, and gauge generators that the canonical packet groups into the five conceptual bundles. Representative SIP cards across protocells, ATP synthase, chemotactic microbes, ring attractors, and collective institutions appear in Appendix I; cross-domain comparability follows from the common grammar, not from identical phenomenology.

Theorem 9.2 (T-SIP: SIP completeness (principal theorem) (T-Qi4-8.2)). *Fix an initial boundary distribution p_0 on $\mathcal{Z}$, and assume:*

- (H1) *Observable completeness. The observable dictionary $\mathcal{O}$ separates points of $\mathcal{Z}$ and the algebra it generates is dense in $C_b(\mathcal{Z})$ under bounded-pointwise convergence.*
- (H2) *Autonomous generator (for the converse). The generator $\mathcal{L}_{\text{GDOS}}$ is autonomous, or equivalently the gauge ambiguity in the converse direction is treated as a time-ordered family of Stinespring–Kraus unitaries rather than a single static unitary.*

Then the full family of operational predictions for every registered observable $O \in \mathcal{O}$ is determined by

$$p_t = \mathcal{U}_{t,0}^{\text{GDOS}}[p_0], \quad \mathcal{U}_{t,0}^{\text{GDOS}} = \mathcal{T} \exp\left(\int_0^t \mathcal{L}_{\text{GDOS}}(\tau; \mathcal{P}(\tau), \eta(\tau)) d\tau\right), \quad (9.3)$$

and $\langle O \rangle_t = \int_{\mathcal{Z}} O(x) p_t(x) \sqrt{\det g} d^n x$. Conversely, two packets are operationally equivalent iff they are related by a diffeomorphic and gauge-covariant pushforward that preserves the observable dictionary and the kill criteria.

Proof. Forward direction. The five fields of the canonical packet are exactly the data required to integrate Eq. (9.3). Field 1 ($\mathcal{Z}$ and μ_g) fixes the domain and integration measure of every observable expectation. Field 2 ($\mathcal{P}(t)$, $\eta(t)$) supplies the boundary couplings entering $\mathcal{L}_{\text{GDOS}}$. Field 3 ($\mathcal{L}_{\text{GDOS}}$) is the GKLS-type generator; by the existence-and-uniqueness theorem for CPTP semigroups under bounded $\mathcal{L}$ (GG-A5-GDOS Theorem T2), it determines p_t uniquely from p_0 for all $t \geq 0$. Field 4 (F_{GFF} , $\mathcal{J}$) supplies the scalar and vector budgets; their role in $\mathcal{L}$ is fixed by Theorem 8.2. Field 5 ($\mathcal{O}$, $\mathfrak{K}$) specifies the experimental contract. Together these determine every prediction $\langle O \rangle_t$ for every $O \in \mathcal{O}$ and $t \geq 0$.

Converse direction. Suppose two canonical packets yield identical predictions for every O , every p_0 , and every $t \geq 0$. *Step 1 (manifolds match).* The observable algebra $\mathcal{A} = \langle \mathcal{O} \rangle$ is a commutative C^* -algebra of bounded functions on $\mathcal{Z}$; by Gelfand–Naimark, equality of all expectation values is equivalent to homeomorphism of the spectra, and smoothness of $\mathcal{L}_{\text{GDOS}}$ upgrades this to a diffeomorphism $\varphi : \mathcal{Z} \rightarrow \tilde{\mathcal{Z}}$. *Step 2 (generators match modulo gauge).* Differentiating the prediction equality at $t = 0$ gives $\mathcal{L}^*[O] = \varphi^* \tilde{\mathcal{L}}^*[\tilde{O}]$; by Stinespring–Kraus uniqueness (Choi, 1975, Kraus, 1983, Stinespring, 1955), two GKLS generators agreeing on a dense subset of the bounded observable algebra are related by a unitary rotation of their jump operators (the Stinespring gauge). *Step 3 (ledgers and observables match).* Once φ and the gauge are fixed, the ledger budgets F_{GFF} and $\tilde{F}_{\text{GFF}}$ coincide on corresponding domains by direct comparison of the aligned observable algebras; Fields 4 and 5 thus match. The reverse implication is the contrapositive: if the two packets are not related by such a pushforward, then there exists an observable, an initial distribution, or a ledger entry on which they disagree, and the packets are not operationally equivalent. ■

Bridge to Part III. With GRIP, GFF, and SIP imported from GG-A6-GRIP as the runtime constitution, the natural next question is what specific four-pillar architecture emerges on the information half of the boundary world. The four Parts that follow answer through the Qi-4 pillars in sequence: Qi-QUA (current locking and the Arisaka equation), Qi-EP (variational principle in three named gauges), Qi-CON (conformality, fractality, $1/f$ two-knee spectra), and Qi-HAL (memory architecture MePMoS-NHT-HAL-BFFT). Each pillar is presented as a dual-layer principal theorem with a named compatibility proposition $\mathfrak{C}_X$; the four pillars together constitute the quantum-information backbone of GDOS.

10 What had to be true first

Nothing in this Part is new, and that is the point of it. A paper that proposes to price information has to say, before it starts, what it already believes about how systems move, what they spend, and how they are written down — because otherwise the price could be made to come out at any figure by adjusting the accounting behind it.

What is on the table now is an engine, a budget, and a standard form. The engine has three moves and two regimes. The budget is a single ledger, and its singleness is not a convenience but a requirement: if exploring, committing and maintaining could each be charged to their own account, there would be nothing to prevent a system from being described as thrifty in one currency while ruinous in another, and no statement about the whole system would mean anything. The standard form carries, in its last field, the conditions under which the description would be judged wrong.

Of the two regimes, one deserves the last word here, because everything this paper adds later grows from it. Descent on a landscape that does not move is a journey with an end; the system rolls, perhaps waits, and eventually reaches somewhere with no exit. That regime explains stability, and it explains why matter is permanent. What it cannot explain is why there was ever anywhere new to go. For that, something has to reorganize the landscape itself — and in cosmology those events are exogenous and unrepeatable. The universe does not re-run its own beginnings.

This is where a universe of atoms and a universe with cells in it part company. A system that remembers can produce such reorganizations from within, and can choose when. That claim is not made in this Part; it is made much later, and it needs the four pillars first. But it is worth seeing now what the imported machinery leaves undone, because the shape of what is missing is the shape of what this paper is for. The engine explains how anything comes to rest. It does not explain how anything came to be able to start again.

The six Parts that follow are the constructive heart of the paper. With the runtime grammar of Part II in place, we now write the four Qi-4 pillars themselves — Qi-QUA (current locking and the Arisaka equation), Qi-EP (variational selector and the unified extremal principle), Qi-CON (conformality, Lévy statistics, fractal supports, and the $1/f$ spectrum), and Qi-HAL (memory architecture from the MePMoS cycle to the holographic attractor lattice). Each pillar is presented as a principal theorem with the L16 dual-layer architecture: an algebraic statement that mathematical physicists can audit, paired with an operational statement that experimenters can falsify.




Figure 4: Architecture overview of the four Qi-4 pillars. Each pillar inherits its constitutional content from the GDOS boundary law of A5 paired with the Gi-4 geometric template of A6 (top), and its bulk-side technical content from the 6D bulk + $U(1)_{Q_i}$ gauge sector (bottom). Horizontal arrows show the implication chain $\text{Qi-QUA} \Rightarrow \text{Qi-EP} \Rightarrow \text{Qi-CON} \Rightarrow \text{Qi-HAL}$. The four named compatibility propositions $\mathcal{C}_{\text{QUA}}, \mathcal{C}_{\text{EP}}, \mathcal{C}_{\text{CON}}, \mathcal{C}_{\text{HAL}}$ (one per pillar) tie the algebraic (T-*.1) and operational (T-*.2) layers within each pillar.

10.1 What this Part returns to the companion it borrows from

The debt recorded above runs one way, and it would leave a false impression to close the Part without recording the other. The geometric companion supplies this paper with an engine, a budget, and a standard form. What it receives back is not a courtesy but a completion, and it is worth naming precisely, because the two papers are halves of one construction and neither is finished alone.

Three things go back. *The budget's informational column becomes physical.* The scalar ledger arrives here as an energetic term plus an effective temperature times a relative entropy, with the second term carrying no price. Qi-QUA prices it: the conversion constant Q_B makes the informational column a quantity in the same units as the first, which is what converts a formal sum into a budget a system can be said to spend. *The cascade acquires an informational reading.* Exploration, commitment and capture are geometric moves in the companion; here they become moves that a system with memory can be said to perform on its own model, which is what allows the same three-move runtime to be recognized in a cell, a nervous system, or an instrument rather than only on a moduli space. *The ground attractor acquires certificates.* The companion establishes where the descent halts; Part VIII establishes what the halting configuration carries — four pillar certificates, proved in that direction, with the converse a theorem on the strata the companion proves admissible, inherited from its one floor, and recorded as open elsewhere.

The relationship is therefore not that one paper is prior and the other applied. The geometric half says how an open system moves and what it spends; the informational half says what the spending is denominated in and what the moving system can come to know. A reader who takes only the first has a runtime with an unpriced column; a reader who takes only the second has a

price without an engine. The two are stated separately because they are proved separately, and they are cited together because neither claim survives on its own.

Part III

Qi-QUA

Current locking, the Arisaka Ratio, the Arisaka Equation, and the Arisaka Constant

Everything that happens at the boundary of an open system costs something. This Part works out what, exactly, and finds that the answer is shorter than anyone had reason to expect.

There are four things a system can spend: energy, electric charge, baryon number, and information. Three of those have had a price for a long time. The fourth has not. Information has been the one quantity in science that everybody bills and nobody prices — invoked whenever a system has to remember, choose or predict, and quietly excused from the accounting every other quantity has to satisfy. The reason is not carelessness. It is that nothing in the standard framework says what a bit ought to cost, so any figure would have been a convention.

What this Part establishes is that the four are not four bills. A single elementary event does not range freely over a four-dimensional space of possible expenditures; it lies on a one-dimensional line inside that space, because all four currents descend through the same channel out of the bulk. Count the events in any one column and you have counted them in all four. The four ledgers are one ledger read in four currencies, and the exchange rate between the energetic column and the informational one is a single constant.


That constant is the Part's second result, and the thing to notice about it is where it comes from. It is not measured. It is the energy scale of the hydrogen atom divided by an integer the ledger fixed several papers upstream, scaled by a coefficient that anomaly cancellation forced. Nothing in the chain was adjusted to make anything come out right, and the Part proves that separately: no new free parameter enters.

Then it can be checked, because life settled on one energy currency and uses it everywhere. The bond that ATP breaks turns out to carry not one of these quanta but two, and the predicted figure lands within one and a half percent of the measured one — inside the experimental error bar. A second check sits in the machine that makes ATP, whose rotor comes in different sizes on different organisms, tuned to the gradient the organism lives on, while the cost of a single click of the rotor stays the same everywhere.

One caution to carry through the Part. The equality is between counted events, not between carriers: it does not say that one electron and one bit are the same object. What it does say, and this is sharp, is that there is no elementary event which carries information while carrying no energy, no charge and no matter. Disembodied information has no entry in this ledger.

The four-pillar sequence begins only after the runtime constitution is complete. The open-system grammar (DGI+GSSB+GA), the scalar ledger (GFF), and the carrier packet (SIP) are all fixed. The first Qi task is therefore event-level locking: what four-sector current ledger is attached to one admissible packet update, and how is that ledger inherited from BI-6? Part III answers that question through the Arisaka Ratio, the Arisaka Constant, and the Arisaka Equation.

Arisaka equation: $J_E/Q_B = J_e = J_B = J_{Qi}$



The four boundary currents lock through one conversion factor Q_B ;
rank-one inflow theorem (T-QUA.1) forces this from BI-6 closure.

Figure 5: The Arisaka equation $J_E/Q_B = J_e = J_B = J_{Qi}$ as the four-current lock of Qi-QUA (Pillar 1 of Qi-4). The four boundary currents — electric J_e , baryonic J_B , informational J_{Qi} , and energetic J_E — are tied through a single energy–information conversion factor $Q_B \approx 0.1555 \text{ eV/qubit}$, the Arisaka constant. The rank-one inflow theorem T-QUA.1 forces this lock from BI-6 closure on the 6D bulk; the empirical match with ATP-hydrolysis chemiosmotic machines at biological temperatures is the operational layer (T-QUA.2).

10.2 Historical background (with citations)

Antecedents: Maxwell (1865), Noether (1918), Bohr (1913), Schrödinger (1944), Landauer (1961), Mitchell (1961), Boyer (1979) (Bohr, 1913, Boyer, 1979, Landauer, 1961, Maxwell, 1865, Mitchell, 1961, Noether, 1918, Schrödinger, 1944); Qi-QUA subsumes these as gauges or limiting cases of Theorem T-QUA.1 $\cup$ T-QUA.2 (the four boundary currents lock through one conversion factor Q_B ; see A1 Sec. 4 and A4 Sec. 3 for full history).

Dual-layer architecture and the compatibility relation $\mathfrak{C}_{\text{QUA}}$. Qi-QUA is, in the L16 language, a dual-layer principal pillar of Qi-4. Its *algebraic layer* (T-QUA.1) is the rank-one inflow theorem (Theorem 12.2): the four boundary sectoral currents J_e , J_B , J_{Qi} , J_E must descend from a single one-dimensional boundary source ray of G_{tot} . Its *operational/physical layer* (T-QUA.2) is the Arisaka equation $J_E/Q_B = J_e = J_B = J_{Qi}$, with the energy–information conversion factor $Q_B \approx 0.1555 \text{ eV/qubit}$ at biological temperatures — one-half the standard free energy of ATP hydrolysis, so that one ATP carries two locked quanta (see the paragraph below and Boyer (1979),

Mitchell (1961)). The two layers are tied together by the named compatibility relation

$$\mathfrak{C}_{\text{QUA}} \equiv \text{T-QUA.1} \implies \text{T-QUA.2} \quad (\text{via the OGN-5 ledger and the boundary projection of } X_4). \quad (10.1)$$

The rank-one inflow at the algebraic level (T-QUA.1) forces the four currents to share one conversion factor; the boundary projection of X_4 via the OGN-5 ledger then identifies that conversion factor with $\mathbf{Q}_B$, the Arisaka constant of T-QUA.2. The physical content of T-QUA.2 is therefore the rank-one structure and the *single* conversion factor, not a numerical accident: in natural packet units the equalities $J_e = J_B = J_{Q_i}$ state that one elementary boundary event deposits the fixed ratio $(1, -1, -2/5, -2/5)$ across the four sectors, $\mathbf{Q}_B$ being the one factor that converts the energetic sector to the others.

The bioenergetic anchor: one ATP carries two quanta. The empirical anchor of $\mathbf{Q}_B$ is the universal energy currency of life. One molecule of ATP carries not one locked quantum but *two*: the standard free energy of hydrolysis of its phosphoanhydride bond is $\Delta G'_{\text{ATP}} = 2 N_A \mathbf{Q}_B \approx 30.0$ kJ/mol, against the measured -30.5 ± 0.5 kJ/mol — an agreement to 1.6%, within the experimental error bar. The factor of two is the two-fold $\mathbf{Q}_B$ -quantum cost of breaking and solvating the bond; ATP is in this sense the seventh ground geometric attractor GGA₇ of the program-wide cascade, with the full derivation in Arisaka (2026, A11-Origin). The comparison is made against this *standard* free energy — an intrinsic property of the molecule fixed by the universal constants of Arisaka (2026, A4-Euclid) — not against the larger *in-vivo* value (≈ 50 kJ/mol), which is not a constant but carries the cell-state mass-action term $RT \ln([ADP][P_i]/[ATP])$; the in-vivo value too is a clean small-integer multiple of $\mathbf{Q}_B$ once the proton stoichiometry is included (next paragraph). It was the recognition that $\mathbf{Q}_B$ is *half* the ATP-hydrolysis free energy, not equal to it, that removed the apparent factor-of-two discrepancy.

The proton stoichiometry and the environment-tuned *c*-ring (forthcoming). Mechanistically, ATP synthase pays for each phosphoanhydride bond with $H^+/\text{ATP} = c/3$ protons, where c is the integer stoichiometry of the membrane rotor and 3 is the number of catalytic sites per turn. Each proton crosses at the proton-motive force Δp , so the in-vivo cost is $(c/3) e \Delta p$. In *E. coli* the rotor is c_{10} (Ballhausen et al., 2009, Jiang et al., 2001), giving $H^+/\text{ATP} = 10/3$ and, at the measured $\Delta p \approx 0.155$ V, a per-proton quantum $e \Delta p \approx \mathbf{Q}_B$ — so the elementary chemiosmotic step *is* the Arisaka quantum, and the in-vivo cost $(10/3) \mathbf{Q}_B \approx 50$ kJ/mol is recovered. The rotor stoichiometry is itself an environment-tuned integer, ranging from c_8 in oxygen-rich mammalian mitochondria (lowest H^+/ATP , highest efficiency) through c_{10} in *E. coli* to c_{13} – c_{15} in low-proton-motive-force alkaliphiles and chloroplasts: the GRIP cascade selects, for each biosphere, the discrete c that matches the available Δp to the universal two-quantum bond. That cross-environment optimization — the integer phenotype of ATP synthesis, and the sense in which ATP-based life is the inevitable bioenergetic attractor of an Earth-like biosphere — is developed in a forthcoming dedicated study.

11 The four ledgers of a generic open system

The appropriate starting point is not a specialized chemical or biological reaction, but the generic bookkeeping structure of an open system. Let S be an open system with operational boundary ∂S . For each sector

$$A \in \{E, e, B, Q\},$$

we introduce a density $n_A(\mathbf{x}, t)$, a flux $\mathbf{j}_A(\mathbf{x}, t)$, and a source term $s_A(\mathbf{x}, t)$ obeying the continuity equation

$$\partial_t n_A + \nabla \cdot \mathbf{j}_A = s_A. \quad (11.1)$$

Equation (11.1) is the open-system version of Noether-style current bookkeeping, with explicit source terms retained because the system is not closed (Arisaka, 2026, A5-GDOS, Noether, 1918). Integrating over a control volume $V \subset \partial S$ gives

$$\frac{d}{dt} \int_V n_A d^3x = - \int_{\partial V} \mathbf{j}_A \cdot d\mathbf{S} + \int_V s_A d^3x. \quad (11.2)$$

This equation is the generic bridge from differential current bookkeeping to experimentally registered budgets.

The four sectors used in Qi-QUA are the following.

Table 8: The four ledgers used by Qi-QUA. Energy, electric charge, baryon number and information run down the first column, each with its symbol, its operational meaning, its natural boundary unit and its interpretation in an open system. The table is the content of the pillar in one page: the claim is not that these four quantities are each conserved, which is standard, but that they are counted in units that a single conversion constant relates — so that the four ledgers are one ledger read in four currencies.

Sector	Symbol	Operational meaning	Natural boundary unit	Interpretation in open systems
Energy	E	Net energetic expenditure or uptake driving a registered step	eV per event, or kJ/mol per event	Work-like or free-energy-like cost required to maintain or update the open system
Electric	e	Charge-transfer or redox-update count	one elementary electric transfer per event	Carrier transport, redox step, or equivalent charge bookkeeping
Baryonic	B	Material/structural rearrangement count	one minimal structural update per event	Rearrangement of matter-bearing degrees of freedom
Information	$Q \equiv Qi$	Registered informational update count	one bit-equivalent update per event	Storage, copying, comparison, or commit operation in the boundary register

Table 8 deliberately mixes language from physics, chemistry, and information theory because Qi-QUA is supposed to operate across all three. The purpose of the table is not to claim that a

neuron, a catalytic cycle, and a gauge theory use identical microscopic carriers. The purpose is to identify the four ledgers that must all be audited if one wants a theory of open systems that scales from one domain to another without changing its constitution.

The essential distinction is between *coefficient-space currents* and *natural-unit currents*. Coefficient-space currents are the signed sector weights inherited directly from BI-6. Natural-unit currents are the magnitudes of registered events after each sector has been expressed in its own operational unit. The Arisaka Ratio lives in the first language. The Arisaka Equation lives in the second. Much confusion is avoided if that distinction is maintained from the outset.

Remark 11.1. The letter Q in the present section stands for the informational current J_{Qi} , written in shortened form once the context is fixed. The notation follows the convention

$$\frac{J_E}{Q_B} = J_e = J_B = J_Q, \quad J_Q \equiv J_{Qi},$$

with the order beginning from J_E so that the presentation remains consistent with the BI-6 ordering.

12 From BI-6 to a one-dimensional boundary source ray

Paper II established that the anomaly-free six-dimensional bulk is organized by a Green–Schwarz-modified Bianchi identity — read in the present program as the GGEP completion of the 1984 GS construction (see the section “BI-6 as the GGEP completion of GS: the geometrization reading” of [Arisaka, 2026, A3-BI6](#)) — of the schematic form

$$dH_3 = X_4, \quad X_4 = a X_E + b X_e + c X_B + d X_{Qi}, \quad (12.1)$$

with the coefficient vector fixed by complete anomaly cancellation to

$$(a, b, c, d) = \left(1, -1, -\frac{2}{5}, -\frac{2}{5}\right) \quad (12.2)$$

in the GG-6 construction ([Arisaka, 2026, A2-GG6](#)). The crucial point for the present paper is that a *single* three-form field strength H_3 feeds all four boundary sectors. This immediately reduces the allowed elementary source vectors on the boundary to a rank-one family.

Remark 12.1 (Normalization convention: reduced-rational BI-6 form). In earlier notes one may also encounter a reciprocal Abelian normalization in which the same locking pattern is written schematically as

$$(1, -1, -5/2, -5/2).$$

The present manuscript adopts the reduced-rational convention

$$(1, -1, -2/5, -2/5)$$

because it is exactly the form whose positive integer clearing yields

$$5 : -5 : -2 : -2,$$

which is the Arisaka Ratio used throughout Paper IV-B, and because it is the same convention that leads directly to the chemical-unit definition

$$Q_B = \frac{2}{5} \frac{E_{\text{Ryd}}}{K}.$$

Thus the choice of $-2/5$ is not an independent physical claim beyond BI-6; it is the normalization convention that keeps the ultraviolet coefficient vector, the integer-cleared Arisaka Ratio, and the chemical-unit map mutually consistent. What must be avoided is mixing the reduced-rational BI-6 convention with the reciprocal integer form $2 : -2 : -5 : -5$, or mixing the reciprocal BI-6 convention with the reduced-rational Arisaka Constant.

To see this concretely, integrate Eq. (12.1) over a compact four-volume W_4 whose boundary projects to a finite operational update on ∂S :

$$\int_{W_4} dH_3 = \int_{W_4} X_4 = a \int_{W_4} X_E + b \int_{W_4} X_e + c \int_{W_4} X_B + d \int_{W_4} X_{Qi}. \quad (12.3)$$

Because the left-hand side depends on one integrated inflow amplitude only, the four right-hand sector contributions cannot vary independently for an elementary synchronized event. There exists a scalar amplitude $\lambda(W_4)$ such that the induced signed boundary source vector satisfies

$$\mathbf{s}' = \lambda \left(1, -1, -\frac{2}{5}, -\frac{2}{5} \right). \quad (12.4)$$

The prime reminds us that Eq. (12.4) is a *coefficient-space* relation: it records the signed sector weights inherited directly from BI-6, before any sector-specific unit conversion has been performed.

Because GFF and SIP are already fixed by Part II, the first Qi theorem can now focus purely on event-level locking rather than on scalar-budget or packet-definition ambiguities.

Theorem 12.2 (Rank-one inflow theorem for elementary Qi-QUA events (T-Qi4-9.3)). *Assume that an elementary open-system update on the boundary is sourced by a single BI-6 inflow channel H_3 . Then the admissible signed sector-current vector for that elementary update lies on the one-dimensional ray*

$$\mathcal{R}_A = \left\{ \lambda \left(1, -1, -\frac{2}{5}, -\frac{2}{5} \right) : \lambda \in \mathbb{R} \right\}. \quad (12.5)$$

Proof. Under the stated assumption the four boundary source components are obtained by projecting one and the same bulk object H_3 through the fixed BI-6 coefficient vector of Eq. (12.2). Hence all four components are proportional to a single scalar amplitude λ , and the set of admissible source vectors is one-dimensional. No second independent direction exists unless a second primitive inflow channel is added, which is excluded by hypothesis. The one-channel hypothesis is not an arbitrary ansatz but a consequence of bulk BI-6 closure: within the GG-6 anomaly-canceled spectrum, the closed 4-form X_4 is unique up to scale, and the Green–Schwarz counter-term forbids any second independent H'_3 on the boundary; see the rank-one criterion (T-BI6-6.1) of Arisaka (2026, A3-BI6), together with its section on BI-6 as the GGEP completion of GS. ■

Theorem 12.2 is the whole geometric content of the first Qi pillar in its most compressed form. It says that the elementary bookkeeping event of the boundary does not span a four-dimensional current space. It spans a one-dimensional ray inside that space. All later equalities of Qi-QUA are just different ways of reading that ray.

13 The Arisaka Ratio and its equivalent representations

We may now state the first of the three central objects.

Definition 13.1 (Arisaka Ratio). The *Arisaka Ratio* is the coefficient-space ray associated with the BL-6 source vector, written in the ordered sector basis (E, e, B, Q) :

$$\text{ArisakaRatio} : \quad J'_E : J'_e : J'_B : J'_{Qi} = 1 : -1 : -\frac{2}{5} : -\frac{2}{5}. \quad (13.1)$$

Equivalently, after clearing denominators by a positive rescaling factor 5,

$$J'_E : J'_e : J'_B : J'_{Qi} = 5 : -5 : -2 : -2. \quad (13.2)$$

The two forms (13.1) and (13.2) are mathematically equivalent because they represent the same ray under a positive scalar rescaling. The first form makes direct contact with the BL-6 coefficient vector. The second form is often easier to read because all entries are integers.

This is also the right place to state the sign convention unambiguously. In the present paper the sign of a coefficient-space current records *orientation*, not whether the corresponding physical quantity is “good” or “bad.” We choose the energy inflow direction as positive. The electric, baryonic, and informational components then carry negative sign in Eq. (13.1) because they are recorded as sink or deposition channels relative to that orientation. When the same bookkeeping is rewritten in natural units, one normally works with magnitudes rather than signed orientations. The Arisaka Equation will therefore be written without minus signs.

Table 9: Equivalent ways of reading the Arisaka Ratio. The first column names the representation, the second gives the formula, and the third says where each form is useful. The three are the same statement: the coefficient form as it falls out of the anomaly normalization, the integer-cleared form that makes the ledger arithmetic visible, and the magnitude form used when only relative sizes matter. They are set out together because the literature of the cohort uses all three and a reader meeting the second without the first can mistake an integer clearing for a different claim.

Representation	Formula	Use
BL-6 coefficient form	$1 : -1 : -2/5 : -2/5$	Closest to the anomaly-cancellation data of GG-6; best for ultraviolet bookkeeping.
Integer-cleared form	$5 : -5 : -2 : -2$	Same ray, but easier to scan and compare across tables.
Magnitude form	$1 : 1 : 2/5 : 2/5$	Obtained by taking absolute values when only event magnitudes are relevant.
Natural-unit form	$J_E/Q_B = J_e = J_B = J_Q$	Best for chemistry, biology, and cognition, where each sector is counted in its own operational unit.

It is important not to overread Definition 13.1. The Arisaka Ratio does *not* claim that every coarse-grained macroscopic current measured in a laboratory or organism must appear in the raw ratio $1 : 1 : 2/5 : 2/5$. What it claims is more precise: the *elementary anomaly-mediated source direction* inherited from the bulk lies on that ray. Coarse-grained systems may contain storage, overhead, idle cycles, and side reactions. Those effects are built from sums of elementary events and therefore belong to later bookkeeping layers. Qi-QUA first identifies the primitive event.

Proposition 13.2 (No independent elementary informational direction). *Within the one-channel hypothesis of Theorem 12.2, there is no elementary boundary event carrying $J'_{Qi} \neq 0$ while $J'_E = J'_e = J'_B = 0$.*

Proof. Any elementary source vector lies on the ray $\mathcal{R}_A$ of Eq. (12.5). On that ray the informational component is proportional to the same scalar amplitude λ that controls all other components. Hence $J'_{Qi} \neq 0$ implies $\lambda \neq 0$, which in turn implies $J'_E \neq 0$, $J'_e \neq 0$, and $J'_B \neq 0$. ■

Proposition 13.2 is the precise framework-level sense in which the present theory rejects “disembodied information” at the elementary ledger level. The statement should not be confused with the weaker claim that every phenomenological reduced model must retain explicit electric and baryonic variables. Reduced models may integrate out those ledgers. The proposition concerns the underlying synchronized event, not every possible reduced description.

Falsification class and kill-criterion. The Arisaka equation $J_E/Q_B = J_e = J_B = J_{Qi}$ is required to hold *as an equality*, within a stated tolerance, at the level of single elementary events in the rank-one inflow channel. The falsification class is therefore concrete and is the only place a hostile experimenter is permitted to look:

- single-molecule kinetics of catalytic engines (ATP synthase rotor step, kinesin step, photosynthetic charge separation, single-enzyme turnover);
- single conformational transitions of memory-bearing biomolecules (CaMKII autophosphorylation event, single ribosomal step, single transcription-factor binding/unbinding);
- single elementary events in microbial electron transport (a single quinone reduction step, a single cytochrome electron transfer).

For each such event, the energy, electron-count, baryon-count, and information-bit bookkeeping must agree to within $\pm 20\%$ — the empirical dispersion of physiological ATP-hydrolysis free energies across normal cellular conditions and the calibration uncertainty in matching the BI-6 ratio to a specific molecular substrate (see Open Problem O5 for the calibration discussion). A systematic violation of the equality outside this tolerance, in a system that is operationally elementary (single event, no integrated-out hidden ledger), refutes Qi-QUA. Coarse-grained or reduced systems are explicitly outside the falsification class: Proposition 16.1 (below) permits $\geq$ inequalities for them because overhead, storage, idle cycles, and side reactions are built from sums of elementary events and belong to later bookkeeping layers. The kill-criterion is therefore: *find one elementary catalytic event whose four-current bookkeeping disagrees with the Arisaka equation outside $\pm 20\%$, and Qi-QUA is refuted.*

14 From GG-6 scales to the Arisaka Constant

The second central object of the section is the fixed energy scale assigned to one locked boundary event in chemical natural units. The derivation uses only quantities already imported from Papers II and III.

The first ingredient is the OGN-5 integer $K = 35$, which in Paper II measured the discrete depth of the six-dimensional throat. The second ingredient is the Rydberg energy

$$E_{\text{Ryd}} = \frac{1}{2}\alpha^2 m_e c^2, \quad (14.1)$$

the characteristic atomic energy scale descending from the previously derived low-energy electromagnetic sector (Arisaka, 2026, A2-GG6, Bohr, 1913). The quantity E_{Ryd} is not introduced here as an arbitrary empirical fit parameter; within the GG sequence it is part of the already imported low-energy output. Numerically $E_{\text{Ryd}} \approx 13.6057 \text{ eV}$.

The next move is to define the minimal geometric step along the chemistry-relevant boundary ledger. The argument for the $1/K$ scaling has two parts.

First, the role of K in the scale-profile throat. In GG-A3-BI6, the OGN-5 integer $K = 35$ is fixed by the Diophantine constraint $K = LJ = M^2 + N^2 + L^2$ with $(M, N, L; J) = (1, 3, 5; 7)$, and it plays a second role as the logarithmic scale depth of the six-dimensional throat: the canonical electroweak-versus-Planck scale identification $M_{\text{lock}} = \frac{1}{2}M_{\text{Pl}} \exp(-K)$ shows that K measures the number of natural log-units between the high-energy (Planck) boundary and the low-energy (electroweak) boundary of the extra direction. This is the same K that we now use to count the discrete depth levels of the atomic boundary.

Second, the linear $1/K$ per-step scaling. A Kaluza–Klein reduction on an interval of radial-coordinate length K (in natural log-units) produces, on the boundary, an evenly spaced ladder of K modes whose total energy span is the atomic scale E_{Ryd} . The per-step quantum is therefore E_{Ryd}/K , exactly the standard linear-spacing result for KK modes on an interval of finite length. An alternative $1/K^2$ scaling (which would occur for a flat compactification with momentum-quantization on a torus) is excluded here because the boundary is the IR boundary, not a flat fiber; an alternative $1/\log K$ scaling (which would occur for an exponential discretization of the throat) is excluded because the OGN-5 ledger fixes the discrete depth by linear, not exponential, increments (Arisaka, 2026, A3-BI6). Hence

$$\delta E_{\text{geom}} \equiv \frac{E_{\text{Ryd}}}{K}. \quad (14.2)$$

Numerically $\delta E_{\text{geom}} \approx 13.6057/35 \approx 0.389 \text{ eV}$.

This quantity is not yet the energy per locked chemical event. It is the per-step geometric energy quantum before the BI-6 coefficient ratio between the energy sector and the baryonic/informational sectors has been applied.

Remark 14.1 (On the apparent double use of the BI-6 ratio $2/5$). The same BI-6 coefficient ratio $|c|/|a| = 2/5$ appears at two places in the construction below: once as the multiplicative factor in $\mathbf{Q}_B = (2/5)\delta E_{\text{geom}}$ (Eq. (14.2) $\rightarrow$ Eq. (14.4)), and once as the reciprocal factor $5/2$ in the chemical-natural-unit map of the baryonic/informational currents (Eq. (15.7) ff.).

These are not two independent uses of the ratio that double-count one another; they are the same single BI-6 ratio applied to a single sector conversion — once as a magnitude (the energy per locked event) and once as the inverse normalization (the chemical unit of the baryonic/informational current), so that the resulting Arisaka equation $J_E/Q_B = J_e = J_B = J_{Qi}$ holds as a pure-number equality with all sectors counted in their own operational units.

The non-trivial empirical content of Q_B is therefore not the numerical value $\approx 0.155 \text{ eV/qubit}$ by itself — which is fixed by the upstream OGN-5 integer K and the Rydberg energy — but the structural assertion that *every* elementary boundary event in the rank-one inflow channel shares the same single conversion factor between energy and information. The falsifier of Qi-QUA (specified in the next section) is the violation of this single-conversion-factor assertion across a class of operationally elementary events, not the disagreement between the numerical value of Q_B and any one calorimetric measurement.

That ratio is precisely the third ingredient. From Eq. (12.2) one reads

$$\frac{|c|}{|a|} = \frac{2}{5}. \quad (14.3)$$

In other words, the magnitude of one baryonic or informational elementary update is $2/5$ of the coefficient-space energy step when the latter is written in the common BI-6 normalization. Therefore the energy cost of one *boundary event counted in the natural baryonic/informational unit* is

$$Q_B = \frac{|c|}{|a|} \delta E_{\text{geom}} = \frac{2}{5} \frac{E_{\text{Ryd}}}{K}. \quad (14.4)$$

Definition 14.2 (Arisaka Constant). The *Arisaka Constant* is the fixed energy scale

$$\text{ArisakaConstant} : \quad Q_B = \frac{2}{5} \frac{E_{\text{Ryd}}}{K}, \quad K = 35. \quad (14.5)$$

It is the energy cost assigned to one ideal locked boundary event when the electric, baryonic, and informational sectors are each counted in one natural unit.

Using $E_{\text{Ryd}} \approx 13.6057 \text{ eV}$ and $K = 35$, one obtains

$$Q_B = \frac{2}{5} \frac{13.6057 \text{ eV}}{35} \approx 0.1555 \text{ eV}. \quad (14.6)$$

The same value may be written in two experimentally familiar ways:

$$Q_B \approx 15.0 \text{ kJ mol}^{-1} \approx 5.82 k_B T \quad \text{at } T = 310 \text{ K}. \quad (14.7)$$

The thermal reading here is not an independent input. The $\text{eV} \leftrightarrow k_B T$ conversion is one face of the thermal clock: the same one-sided (positive-frequency) spectrum that endows the informational carrier with its derived quarter-turn returns, in KMS form, the bath temperature (T-P5-B1 of [Arisaka, 2026, P5-CPT](#); candidate reading).

Table 10: Numerical forms of the Arisaka constant. The first column names the form, the second gives the value and the third its interpretation. The four rows are one number — the definition from the geometry, then the same quantity in atomic, chemical and thermal units, the last evaluated at 310 K. The chemical row is the one that meets experiment: it is what makes the ATP anchor a comparison rather than an analogy, and the thermal row is what shows the constant is not a disguised $k_B T$.

Form	Value	Interpretation
Definition from GG-6	$Q_B = \frac{2}{5} \frac{E_{\text{Ryd}}}{K}$	No new continuous parameter beyond the imported GG-6 data
Atomic-energy form	$Q_B \approx 0.1555 \text{ eV}$	Energy scale of one locked boundary event
Chemical form	$Q_B \approx 15.0 \text{ kJ mol}^{-1}$	Convenient for chemistry, catalysis, and biophysics
Thermal form (310 K)	$Q_B \approx 5.82 k_B T$	Useful for room-temperature or physiological open systems

The significance of Definition 14.2 is structural. The paper is not merely introducing a convenient conversion factor. It is asserting that the energy cost per ideal locked event is not free to float continuously from system to system. What may vary from system to system is *how many* locked events are required, how much additional overhead is paid, and which geometric attractor is realized. The per-event baseline Q_B is fixed by the imported GG-6 data.

Theorem 14.3 (No new continuous parameter in Qi-QUA(T-Qi4-9.8)). *Within the GG-6/GDOS import structure of Paper IV-B, the constant Q_B introduces no new continuous free parameter.*

Proof. Equation (14.5) expresses Q_B in terms of the already imported quantities E_{Ryd} and K . The Rydberg scale is part of the low-energy electromagnetic output inherited from the previous papers, while $K = 35$ is one of the fixed OGN-5 integers. Therefore Q_B is derived data, not an adjustable input. ■

15 Natural-unit reduction and the Arisaka Equation

We now convert the coefficient-space statement into the natural-unit statement appropriate for open systems in chemistry, biology, and cognition.

The single most important conceptual point is that the Arisaka Equation is *not* obtained from the Arisaka Ratio by merely dropping minus signs. One must first perform a sector-dependent natural-unit conversion. Stated in the exact order used in the present manuscript, the derivation is

$$\text{BI-6 normalization:} \quad (a, b, c, d) = \left(1, -1, -\frac{2}{5}, -\frac{2}{5}\right), \quad (15.1)$$

$$\text{Arisaka Ratio:} \quad J'_E : J'_e : J'_B : J'_{Qi} = 1 : -1 : -\frac{2}{5} : -\frac{2}{5} = 5 : -5 : -2 : -2, \quad (15.2)$$

$$\text{chemical-unit map:} \quad \frac{J_E}{Q_B} : J_e : J_B : J_Q = |J'_E| : |J'_e| : \frac{5}{2}|J'_B| : \frac{5}{2}|J'_{Qi}|, \quad (15.3)$$

$$\text{Arisaka Equation:} \quad \frac{J_E}{Q_B} = J_e = J_B = J_Q. \quad (15.4)$$

Equation (15.1) is the ultraviolet normalization choice. Equation (15.2) is the coefficient-space statement inherited from BI-6. Equation (15.3) is the indispensable conversion from raw sector coefficients to chemical natural units. Only after that conversion does Eq. (15.4) follow.

To make the same point in slower arithmetic, suppose an elementary synchronized event carries coefficient-space magnitudes

$$(|J'_E|, |J'_e|, |J'_B|, |J'_{Qi}|) = r \left(1, 1, \frac{2}{5}, \frac{2}{5} \right), \quad r \geq 0, \quad (15.5)$$

for some event-rate amplitude r . In other words, the raw coefficient-space bookkeeping says only that the baryonic and informational components are smaller than the electric and energy components by the factor $2/5$ in the BI-6 normalization.

The next step is to define the *chemical natural units*. We choose the electric sector as the reference event count and define

$$J_e \equiv |J'_e|, \quad (15.6)$$

$$J_B \equiv \frac{5}{2} |J'_B|, \quad (15.7)$$

$$J_Q \equiv \frac{5}{2} |J'_{Qi}|, \quad J_Q \equiv J_{Qi}, \quad (15.8)$$

$$\frac{J_E}{Q_B} \equiv |J'_E|, \quad \text{equivalently} \quad J_E \equiv Q_B |J'_E|. \quad (15.9)$$

These definitions do exactly one thing. They convert the four raw BI-6 amplitudes into *event counts* in the natural unit of each sector:

- the electric sector is counted one electric-transfer event at a time;
- the baryonic sector is counted one baryonic-update event at a time, which requires multiplication by the reciprocal factor $5/2$ because its coefficient-space magnitude is $2/5$ of the electric one;
- the informational sector is counted one informational-update event at a time, again requiring the reciprocal factor $5/2$;
- the energy sector is counted in units of the fixed event energy Q_B , so that J_E/Q_B is the corresponding event count.

Substituting Eq. (15.5) into Eqs. (15.6)–(15.9) yields

$$J_e = r, \quad (15.10)$$

$$J_B = \frac{5}{2} \left(\frac{2}{5} r \right) = r, \quad (15.11)$$

$$J_Q = \frac{5}{2} \left(\frac{2}{5} r \right) = r, \quad (15.12)$$

$$\frac{J_E}{Q_B} = r. \quad (15.13)$$

Therefore

$$\frac{J_E}{Q_B} = J_e = J_B = J_Q = r. \quad (15.14)$$

We may now state the third central object.

Definition 15.1 (Arisaka Equation). The *Arisaka Equation* is the natural-unit form of the Qi-QUA locked current ray:

$$\text{ArisakaEquation} : \quad \frac{J_E}{Q_B} = J_e = J_B = J_Q, \quad J_Q \equiv J_{Qi}. \quad (15.15)$$

Remark 15.2 (Why $J_e = J_B = J_Q$ is not a raw microscopic identity). The equality in Eq. (15.15) does *not* mean that one electron, one baryon, and one bit are literally the same microscopic carrier. In coefficient-space units the correct statement is instead

$$|J'_E| : |J'_e| : |J'_B| : |J'_{Qi}| = 1 : 1 : \frac{2}{5} : \frac{2}{5}.$$

The equalities $J_e = J_B = J_Q$ appear only after the baryonic and informational sectors are measured in their own natural packet units, that is, after the reciprocal 5/2 conversion has been applied. The Arisaka Equation is therefore an equality of *normalized event counts*, not an identity among raw microscopic carriers.

This is the formal bridge from coefficient-space BI-6 data to chemistry/open-system event counting, and it is the mathematical reason Qi-QUA belongs after GFF and SIP rather than before them.

Theorem 15.3 (Natural-unit equivalence theorem (T-Qi4-9.11)). *For an elementary boundary event sourced by one BI-6 inflow channel, the coefficient-space Arisaka Ratio of Eqs. (13.1)–(13.2) is equivalent to the natural-unit Arisaka Equation (15.15).*

Proof. The coefficient-space ray gives Eq. (15.5). Applying the chemical-unit map (15.3), equivalently the definitions (15.6)–(15.9), yields the explicit equalities (15.10)–(15.13). These show that all four normalized event counts are equal to the same amplitude r . Hence Eq. (15.15) follows.

Conversely, if Eq. (15.15) holds and the chemical-unit map is inverted, then

$$|J'_e| = \frac{J_E}{Q_B}, \quad |J'_B| = \frac{2}{5}J_B, \quad |J'_{Qi}| = \frac{2}{5}J_Q,$$

so the coefficient-space magnitudes necessarily lie on the ray $(1, 1, 2/5, 2/5)$. Reintroducing the BI-6 orientation signs reproduces Eqs. (13.1)–(13.2). The two statements are therefore equivalent once the sector-specific unit map has been fixed. ■

Theorem 15.3 is the formal bridge between particle-physics notation and chemistry-natural-unit notation. In the ultraviolet language one sees the signed coefficient vector inherited from anomaly cancellation. In the chemistry-natural-unit language one sees the equality of per-event ledgers. Both are the same statement, but only after the chemical-unit map is inserted explicitly.

Table 11: Explicit four-line derivation chain of Qi-QUA. Each row is one step: the anomaly normalization that fixes the coefficients, the Arisaka Ratio that follows from it, the map into chemical units, and the Arisaka equation itself. The second column carries the formula and the third what the step means. The chain is short deliberately — there is no intermediate modeling step in it, so a reader who accepts the first row and the arithmetic has accepted the pillar.

Line	Formula	Meaning
BI-6 normalization	$(a, b, c, d) = (1, -1, -2/5, -2/5)$	Ultraviolet reduced-rational normalization chosen so the later current ledger remains compatible with GG-6 and OGN-5.
Arisaka Ratio	$J'_E : J'_e : J'_B : J'_{Q_i} = 1 : -1 : -2/5 : -2/5 = 5 : -5 : -2 : -2$	Coefficient-space source ray inherited directly from BI-6.
Chemical-unit map	$J_e = J'_e , J_B = (5/2) J'_B , J_Q = (5/2) J'_{Q_i} , J_E = Q_B J'_E $	Sector-dependent unit conversion from raw BI-6 amplitudes to event counts and energy units.
Arisaka Equation	$J_E/Q_B = J_e = J_B = J_Q$	Equality of normalized event counts for one ideal locked event.

16 Ideal equality, coarse-grained inequality, and overhead

The Arisaka Equation is exact for the *ideal synchronized elementary event*. Real chemistry, metabolism, and cognition, however, are not made of one event only. They consist of long histories containing useful steps, housekeeping steps, idle cycles, redundancy, proofreading, error correction, and heat leakage. It is therefore essential to state clearly how Qi-QUA behaves under coarse-graining.

Let a macroscopic process be decomposed into N_{lock} ideal locked events plus additional overhead events that increase the energy, electric, or baryonic budgets without increasing the registered informational count. Then the integrated budgets over the whole process can be written schematically as

$$\mathcal{E} = Q_B N_{\text{lock}} + \Delta_E^{\text{over}}, \quad (16.1)$$

$$\mathcal{N}_e = N_{\text{lock}} + \Delta_e^{\text{over}}, \quad (16.2)$$

$$\mathcal{N}_B = N_{\text{lock}} + \Delta_B^{\text{over}}, \quad (16.3)$$

$$\mathcal{N}_Q = N_{\text{lock}}, \quad (16.4)$$

with nonnegative overhead terms $\Delta_A^{\text{over}} \geq 0$. Under this monotone-overhead assumption one immediately obtains

$$\mathcal{E} \geq Q_B \mathcal{N}_Q, \quad \mathcal{N}_e \geq \mathcal{N}_Q, \quad \mathcal{N}_B \geq \mathcal{N}_Q. \quad (16.5)$$

Proposition 16.1 (Coarse-grained lower-bound form of Qi-QUA). *If a macroscopic process is composed of ideal locked events plus nonnegative overhead in the energy, electric, and baryonic sectors, then Eq. (15.15) descends to the lower-bound inequalities (16.5).*

Proof. Substitute Eqs. (16.1)–(16.4) into each difference $\mathcal{E} - Q_B \mathcal{N}_Q$, $\mathcal{N}_e - \mathcal{N}_Q$, and $\mathcal{N}_B - \mathcal{N}_Q$. The differences are precisely the nonnegative overhead terms. ■

Proposition 16.1 is indispensable for realistic applications. Without it one would be tempted either to overstate Qi-QUA as a literal equality for every coarse-grained measured flux, or to dismiss it whenever a real molecular machine or nervous system displays inefficiency. The correct reading is simpler. Qi-QUA fixes the primitive event and therefore also the ideal lower bound. Real systems may and usually do spend more.

Table 12: Ideal and coarse-grained forms of Qi-QUA. Three levels of description run down the first column — a single synchronized elementary event, a macroscopic process carrying overhead, and a phenomenological reduced model — with the mathematical form and the interpretation alongside. The distinction matters for falsification: the equality is asserted at the first level only, and the second and third rows carry inequalities, so a coarse-grained measurement that misses the ideal value is not a refutation and the table says so.

Level of description	Mathematical form	Interpretation
Elementary synchronized event	$J_E/Q_B = J_e = J_B = J_Q$	Exact equality on the ideal locked ray
Macroscopic process with overhead	$\mathcal{E} \geq Q_B \mathcal{N}_Q, \mathcal{N}_e \geq \mathcal{N}_Q,$ $\mathcal{N}_B \geq \mathcal{N}_Q$	Lower-bound form after aggregation and inefficiency
Phenomenological reduced model	Some ledgers may be hidden	Acceptable only if hidden ledgers can in principle be restored in the full bookkeeping

This distinction between exact equality and coarse-grained lower bound is especially important for later biological applications. A protocell, an ATP synthase complex, a genetic proofreading network, or a cortical microcircuit can all spend substantially more than the ideal floor. What Qi-QUA claims is that the floor is not arbitrary. It is fixed by the same imported geometry that fixed the coefficient-space ray.

17 Relation to Landauer and to classical nonequilibrium thermodynamics

One immediate use of the Arizaka Constant is to compare the Qi-QUA event floor with the classical lower bound for logically irreversible erasure. Landauer’s bound states that erasing one bit of information in thermal equilibrium at temperature T requires at least

$$E_L = k_B T \ln 2 \quad (17.1)$$

of dissipated energy (Landauer, 1961). At $T = 310$ K,

$$E_L \approx 0.0185 \text{ eV}. \quad (17.2)$$

Comparing with Eq. (14.6) gives

$$\left. \frac{Q_B}{k_B T \ln 2} \right|_{T=310 \text{ K}} \approx 8.40. \quad (17.3)$$

The meaning of Eq. (17.3) must be stated carefully. Qi-QUA does *not* refute Landauer’s principle. The two statements answer different questions. Landauer asks for the minimum dissipation associated with an idealized logically irreversible operation. Qi-QUA asks for the minimum cost of a *locked open-system boundary event* in which energy, charge, baryonic structure, and information are audited simultaneously. Because the latter problem is richer, its lower bound can legitimately exceed the former.

Historically, the need to distinguish these levels was already implicit in nonequilibrium thermodynamics. Onsager’s reciprocity, Prigogine’s irreversible thermodynamics, Jarzynski’s equality, Crooks’ fluctuation relation, and modern stochastic thermodynamics all teach the same lesson: once a system is open and fluctuating, one must specify which ledger is being optimized, inferred, or dissipated (Crooks, 1999, Jarzynski, 1997, Onsager, 1931, Prigogine, 1967, Seifert, 2012). Qi-QUA adds one extra layer to that tradition. It does not only track entropy production or work. It also insists that charge, material rearrangement, and information update belong to the same primitive event ledger.

Table 13: Landauer’s bound and Qi-QUA answer different lower-bound questions. The two principles occupy the two rows; the columns give the question each addresses, the bound it delivers and the ledger it audits. The table exists to forestall a natural objection. Qi-QUA’s quantum exceeds $k_B T \ln 2$, which looks like a violation until one notices that Landauer audits one fiber — the thermal cost of erasure — while Qi-QUA audits the cost of writing on both fibers at once. They are not competing bounds on the same quantity.

Principle	Question addressed	Lower bound	Ledger audited
Landauer	What is the minimum dissipation for idealized logical erasure of one bit?	$k_B T \ln 2$	Information plus thermal bath
Qi-QUA	What is the minimum cost of one locked boundary event?	$Q_B = (2/5)E_{\text{Ryd}}/K$	Energy, electric, baryonic, and informational ledgers together

The payoff of Table 13 is conceptual economy. It prevents later sections from confusing thermodynamic erasure, biochemical update, and observer-level measurement. These processes are related, but they are not identical. Qi-QUA provides the ledger common to all of them once the full open-system event is audited.

Worked examples. CO activation and hydrogenation in the interstellar medium (Fuchs et al., 2009), ATP synthase in the c_{10} ring of *E. coli* (Ballhausen et al., 2009, Jiang et al., 2001, Yoshida et al., 2001), and DNA structure/function/dynamics (Benham, 1979, Fuller, 1971, Wang, 1979, Watson and Crick, 1953) are recorded as packet-normalized audits in Appendix E; each example tabulates N_{lock} and the corresponding $J_E = N_{\text{lock}}Q_B$.




Figure 6: The anatomy of the two fibers. (a) The compact phase circle: every path closes, so everything a connection does there is holonomy — quantized (the dots are the ledger-rational marks), gauge-protected, transported by a group. (b) The scale interval: no path closes, so nothing is protected; transport composes only forward, a semigroup, and the descent visits each scale exactly once. This is why a stable bit is holonomy-class data on the circle while erasure is a one-way write along the interval, why Landauer’s bound prices only the second of those, and why the two audits are not competing bounds on one quantity. Carried from the holonomy white paper, where the anatomy is the subject rather than a tool.

18 Information and the two fibers: memory on the circle, erasure on the interval

The Landauer comparison above gains a geometric floor from the cohort’s recent campaign (GG-W1-Symmetry; Arisaka, 2026, A5-GDOS; Arisaka, 2026, A8-QM; Arisaka, 2026, P5-CPT), and this paper — the information pillar — is where the reading belongs. The instrument is one exact statement: *holonomy requires closed loops, and an interval has none*. The bulk’s compact circle is nothing but closed loops — everything on it is holonomy, quantized, gauge-invariant, protected; the radial interval can hold no holonomy at all — no protected content, transport composing only forward, a semigroup. Read informationally: *memory lives on the circle; erasure pays the interval*. A bit that persists — a winding, a phase, a parity, a charge — is holonomy-class data on the fiber that remembers, which is why stable records exist at all in an open universe: they are stored in quantities no boundary operation can continuously deform. Erasure — the logically irreversible reset that Landauer priced — is an interval-side write: irreversible not by engineering failure but by geometry, because the fiber it is written along has no inverse transport. Landauer’s bound (17.1) is, in this reading, *the interval’s toll*: the minimum the semigroup charges for one one-way write at temperature T .

Table 13 then acquires a second caption. Landauer audits *one fiber*: the thermal cost of the interval-side write, information against the bath. Qi-QUA audits *both*: a locked boundary event deposits its result in circle-side ledgers (charge, baryonic structure, the informational unit of $U(1)_{Q_i}$ — one rearrangement, one unit) *and* pays the interval-side toll of registration, which is why its floor Q_B legitimately exceeds $k_B T \ln 2$: the richer bound is the price of writing on both fibers at once. The Arisaka Constant is thereby read as an *exchange rate between the circle’s two currencies* —

energy per winding on one face, energy per qubit on the other — fixed not by a bath temperature but by the ledger integers, which is the structural reason it is temperature-independent where Landauer's bound is not. And the measurement chain closes at its companion homes: the record whose writing the quantum companion located on the interval (Arizaka, 2026, A8-QM) is the same record whose cost this section prices and whose persistence the circle guarantees — one event, three audits: *the circle stores it, the interval seals it, and Qi-QUA bills it*. This reading is offered at the register of the campaign — an organization of proved structures, not a new theorem — and every clause in it carries its parent's grade.

19 What it means to put a price on a bit

A price is a strong thing to attach to a quantity. It means the quantity can be compared with others, spent, saved, and run out of. It means somebody can be shown to have got the sum wrong.

What this Part has done, said without symbols, is to reduce four separate accounts to one. A system of the kind considered here was being described as though it kept four books — one for charge, one for the winding of its internal circle, one for information, one for energy — and nothing in the description prevented those books from being kept in different units by four different clerks who never spoke. They are now one book. Three of the columns are locked to each other outright, and the fourth is joined to them by a single number that is the same number everywhere, fixed once by the geometry and never adjusted since. That last clause is what makes the result a price rather than a conversion table. A table has entries, and entries can be revised one at a time by whoever is losing the argument. A price has one value, and everybody is held to it.

For a century, energy has had that status in science and information has not. Information was always the thing you could invoke when the physics ran out — the system "chose", the molecule "recognized", the network "learned" — and no bill ever came. That was not dishonesty; it was the absence of a rate. Without an exchange rate between energy and information, no statement about how much a memory costs is checkable, and so no such statement is really a physical statement at all.

What this Part supplies is the rate. And the thing worth appreciating is where it comes from, because it would have been much easier, and much weaker, to measure it. A number fitted to biochemistry would explain biochemistry and nothing else. This one is assembled out of the energy scale of the hydrogen atom and an integer that a Diophantine closure fixed for reasons that had nothing to do with life, several papers before life was mentioned. That it then lands within one and a half per cent of the free energy of the bond every living thing on this planet uses to pay for everything is either a coincidence or the most interesting sentence in the paper.

It is worth being exact about what a laboratory could do to end this, because it is not what one might guess. The vulnerable claim is not the number. A calorimeter that disagreed with the value would be arguing with the ledger upstream, not with this Part. What is at risk is the assertion that *every* elementary event of this channel shares the *same* single conversion factor — that the enzyme, the electron transfer, the molecular motor and the memory write are all being billed at one rate.

One elementary catalytic event whose four-way accounting misses by more than a fifth would finish it.

It is also worth seeing what the single budget makes possible, because it is the reason this Part comes first rather than somewhere else. A ledger says what a system may spend. It does not say what the system will do. Between all the ways of spending a fixed budget something has to choose one, and until the budget was single there was nothing for such a rule to act on — a rule that minimized one currency while another ran free would not be a rule about the system at all, only about the bookkeeping. Now that the four columns are one, the question of selection is finally a question with a subject. That is the next Part's business.

There is a last thought worth carrying forward. If information really does have a price fixed by geometry, then the familiar picture in which physics describes the world and information theory describes our description of it has the relationship backwards. On this account there is one ledger, it has four columns, and the column we have been treating as bookkeeping is as physical as the other three. Nothing that remembers is getting anything for free.

Part IV

Qi-EP

One projected functional in three gauges: PLT, MEPP, and FEP

Three sciences already know which way an open system goes, and they do not appear to agree.

Optics and mechanics say a system takes the path that makes a certain time or action stationary — the principle that sends light along the quickest route through glass. Nonequilibrium thermodynamics says a driven system settles into the regime that extremizes entropy production, and has spent decades arguing internally about whether that means the most or the least. Theoretical neuroscience says a brain acts so as to minimize the difference between what it predicted and what arrived. Three fields, three foundational principles, three literatures that seldom cite one another, and no reason offered why a cell obeying all three at once should not be pulled in three directions.

This Part shows they are one principle. Not analogous, not complementary: one functional, projected from the six-dimensional action onto the boundary, read three times with a different quantity held primary each time. Hold the path primary and the stationarity condition returns the optical law. Hold the currents primary and it returns the flux–force relation of irreversible thermodynamics. Hold the beliefs primary and it returns the inference rule the brain sciences use. The three are gauges of one object, in the way that Celsius and Fahrenheit are readings of one temperature, and the argument is not a verbal identification but a derivation in each case.

The old quarrel gets settled on the way through. Maximum entropy production and Prigogine’s minimum entropy production, which have been treated as rival slogans for seventy years, turn out to be conjugate readings of the same stationary point under dual conventions about what is held fixed — the first selecting the transport law, the second the state. They were never in competition; they were answering different questions of one object. (A third and distinct statement, Onsager’s least dissipation, is the flux-controlled Legendre transform of the same stationary point, and is derived alongside them; it is easy, and common, to conflate the two minimum principles, and this Part keeps them apart.)

What this buys, beyond tidiness, is that a single system can be described in all three languages without contradiction. A cell following a chemical gradient traces a path, maintains a set of metabolic fluxes, and updates an internal estimate of where the food is. Three descriptions, three vocabularies, one budget underneath — and the Part is where that budget is written down.

With event-level bookkeeping fixed by Qi-QUA, the next logical question is selection. A packet can obey the right ledger and still admit many admissible histories. Part IV therefore derives the selector layer of the Qi program: once the budgeted event cone is known, which trajectory, branch, or inferential state is actually extremized?

Every extremal descent in this Part runs one way: admissible histories fall to the selected trajectory, current, or posterior and do not climb back. That one-way character is consumed here, not created. Per the constitutional custody chain, AX-3 owns the arrow — the one-way reading of

the bulk-to-boundary dictionary is a clause of the constitution (the AX-3 synthesis of [Arizaka, 2026, A1-Constitution](#)) — with the canonical law-level derivation in Part XIV of the CPT member paper over the exactly T -symmetric bulk ([Arizaka, 2026, P5-CPT](#)), and the dynamical implementation in the GDOS boundary law ([Arizaka, 2026, A5-GDOS](#)) and its GG-A6-GRIP runtime pipeline ([Arizaka, 2026, A6-GRIP](#)). The Qi-EP selector, and the whole Qi-4 runtime it drives, is a downstream consumer of that single constitutional arrow, not a private source of it.

The Part now runs in two movements. The first (Sections 20–24) gives the three principles their own histories at review depth — what each says exactly, where each is proved, where each fails, and who located the failure — because a derivation is worth only as much as the precision of what it derives, and because the reader this Part most wants to reach arrives owing allegiance to one of the three literatures and suspicion to the other two. The second movement (Section 25 onward) is the derivation itself: the constitutional statement and the contract it must meet, the projected functional, the three gauges each obtained as a theorem with its regime attached, the failure theorem that bounds all three, the laboratory record, and the unified theorem that closes the Part.

20 Overview of PLT, MEPP, and FEP

Three principles, from three sciences, answer the same question — *which way does an open system go?* — and each answers it with an extremum.

The oldest is the principle of least action, with its optical face, the principle of least time (PLA/PLT): a path is realized when a certain time or action is stationary against all nearby paths. It is the most successful variational idea in science — geometric optics, analytical mechanics, general relativity, and the Standard Model are all written in its language — and it is silent, as classically stated, about systems that are open, driven, and noisy.

The second lives in nonequilibrium thermodynamics and has spent seventy years arguing with itself. Near equilibrium, Prigogine’s theorem of 1945 says a gently driven system settles into the steady state of *minimum* entropy production — one of the few statements in the subject that is actually proved. A rival tradition, running from the theory of plastic materials to the modeling of planetary climates, says driven systems arrange themselves to produce entropy *fastest* (MEPP) — a rule that, fed nothing but sunlight and an energy balance, reproduces the equator-to-pole temperature contrast on Earth, Mars, and Titan.

The third is the youngest and the most ambitious. The free energy principle (FEP) of theoretical neuroscience holds that anything which persists behaves so as to make its sensory input unsurprising — that a brain, and in the strongest reading any living thing, is a prediction machine acting to minimize a variational free energy.

A physicist meeting the three side by side notices two things at once. The first is that they cannot all be fundamental as stated: a living cell is a near-equilibrium chemical network, a driven thermodynamic engine, and a predictive controller *simultaneously*, and if the three principles were independent laws the cell would be receiving three sets of instructions with no account of why they do not conflict. The second is that none of the three has ever been derived. Each is either a theorem

confined to a regime far narrower than its use, or a postulate whose attempted derivations failed at published, identifiable steps. The three failures, moreover, have the same anatomy. No principle announces its own domain of validity — each was applied outside its regime for years before its own community mapped the boundary. Each leaves a load-bearing choice to the person applying it — which entropy production to extremize, where to draw the boundary between system and world. And each general derivation that was seriously attempted broke at a specific step that its own literature has since located. When three literatures that seldom cite one another fail in the same three ways, the failure stops being three pieces of bad luck and becomes a symptom: evidence of one missing structure underneath all three.

This Part supplies the missing structure and proves it is one. The claim, stated once here and then earned: the three principles are three *gauges* of a single functional, projected from the six-dimensional bulk action onto the boundary through OGGEp and GDOS — and the regimes, the quarrels, and the failures of all three follow from the same projection as theorems. Sections 21 through 24 first give each principle its own history at review depth — what it says exactly, where it is proved, where it fails, and who located the failure — because a derivation is worth only as much as the precision of what it derives. The constitutional statement (Section 25) then fixes the contract a genuine derivation must meet, the projected functional is constructed (Section 26), and the three gauges are derived in turn (Sections 27, 28, 29), each with its regime theorem and its laboratory record (Sections 30, 31). The extended treatment — the complete proofs at review depth and the full critical literatures — is the companion supplement (Arisaka, 2026, S1-EP); the present Part is self-contained.

21 The principle of least action and least time: history, applications, and limitations

21.1 From Fermat to Hamilton

The action tradition begins with light. In 1662 Pierre de Fermat derived the sine law of refraction from a single postulate — that light traveling between two points takes the path of least time — and announced the result, with some astonishment at its success, in his correspondence on refraction (Fermat, 1662). The derivation was controversial from the day it appeared: Fermat’s critics in the Cartesian school objected that light has no purposes, and the objection — *why should nature optimize?* — has shadowed the principle in every one of its later forms.

The mechanical version arrived in 1744 twice. Maupertuis proposed that in all natural processes a quantity he called action — mass times speed times distance — is a minimum, and read the economy theologically, as evidence of design (Maupertuis, 1744). In the same year Euler, in the appendix to his treatise on the calculus of variations, gave the statement its first rigorous form: for a single particle at fixed energy, the trajectory extremizes $\int v \, ds$ (Euler, 1744). Lagrange removed both the theology and the restriction, building the whole of mechanics on the variational calculus in the *Mécanique analytique* (Lagrange, 1788); and Hamilton, in 1834, gave the modern form — the realized motion renders the time integral of $T - V$ stationary — and, in the same stroke,




Figure 7: The genealogy of the three extremal principles, drawn as three columns that never cite one another. Left, the action tradition: from Fermat’s least time through Maupertuis, Euler, Lagrange, and Hamilton to Feynman’s stationary phase and the Onsager–Machlup stochastic action. Center, the entropy-production tradition: from the least-heat and least-dissipation theorems of Kirchhoff, Helmholtz, and Rayleigh through Onsager and Prigogine to Ziegler’s maximum, Paltridge’s climates, and the claimed general derivation that was refuted and conceded. Right, the inference tradition: from Helmholtz’s unconscious inference through cybernetic regulation and the Helmholtz machine to predictive coding, the free energy principle, its physics phase, and the critiques that forced its retreat. Each column ends unresolved in its own literature; the convergence at the foot — one projected functional whose three gauges they are — is the claim this Part derives, and the dates above it are why the derivation is worth having.

exhibited the deep identity of optics and mechanics: one characteristic function serves both, light rays and particle paths being two readings of a single variational geometry (Hamilton, 1834). Jacobi sharpened the fixed-energy form soon after: at energy E the orbit of a particle in a potential V extremizes the abbreviated action

$$W = \int \sqrt{2m(E - V)} \, ds, \quad (21.1)$$

which is Fermat's principle with $\sqrt{2m(E - V)}$ playing the role of the refractive index — mechanics as the optics of a medium whose density is set by the potential. That identity is not a curiosity; it is the template for the trajectory gauge below, where the projection supplies the effective index for an arbitrary open system. By the middle of the nineteenth century the principle had become what it remains today: the standard form in which fundamental physical law is written.

21.2 The twentieth century: why “least,” and what noise does

Two twentieth-century results changed the principle's meaning rather than its use. The first explains *why* extremal paths are realized at all. In Feynman's path-integral formulation of quantum mechanics, every path contributes an amplitude of unit magnitude and phase $S/\hbar$; paths near a stationary point of the action add coherently while all others cancel, so the classical trajectory is not obeyed but *survives* — it is the path whose neighbors agree with it (Feynman, 1948). The variational principle is thereby demoted from law to limit, and the demotion is clarifying: it says the principle should be expected to hold exactly where interference is sharp, and to blur where it is not — a statement about regimes, of exactly the kind the other two traditions of this Part never had.

The second result extends the principle to noise. Onsager and Machlup showed that for a linear dissipative system driven by thermal fluctuations, the probability of a trajectory is governed by an exponent of action form,

$$p[\dot{\alpha}] \propto \exp \left\{ \frac{1}{2k_B} \int (\dot{S} - \Phi - \Psi) \, dt \right\}, \quad (21.2)$$

with Φ and Ψ the flux- and force-space dissipation functions — so that the *most probable* path extremizes a functional built from the dissipation, and least dissipation becomes literal: the likeliest history is the least-dissipative one (Onsager and Machlup, 1953). This is the seed of the modern large-deviation viewpoint, and it is the hinge on which this Part will turn: in open systems the action principle survives as a statement about the concentration of path probability, exact in the small-noise limit, and the “time” it extremizes is an effective, geometry-laden time rather than a bare clock reading.

21.3 Applications

The scorecard of the action tradition is unmatched in physics. Geometric optics is Fermat's principle applied; analytical mechanics is Hamilton's; general relativity takes the Einstein–Hilbert action as its definition, and the Standard Model is specified, in its entirety, by a Lagrangian density. Where the tradition is weakest is exactly where this Part operates: open systems. A cell swimming

up a gradient traces a path as surely as a photon crossing a lens, and the question whether any least-time statement governs that path is not answered by any of the classical results above.

21.4 Limitations

Five limitations bound the classical principle, and each is load-bearing for what follows.

First, “least” is a misnomer. The realized path renders the action stationary, not minimal; beyond a kinetic focus the classical path is a saddle point, and elementary examples abound in which no minimum exists (Gray and Taylor, 2007). The correct universal statement is stationarity — which matters here because the projected functional of this Part is likewise stationary, not minimized, and its three gauges inherit exactly that character.

Second, dissipation resists the Lagrangian frame. Bauer proved in 1931 that the equations of motion of a linearly damped system admit no variational formulation in the original variables (Bauer, 1931): friction cannot simply be appended to a Lagrangian. Rayleigh’s dissipation function is bolted onto the Euler–Lagrange equations from outside, not derived from a stationarity principle. For two centuries the action tradition therefore stopped, in principle, at the boundary of every open system.

Third, the stochastic extension is a probability statement, not a dynamics. The Onsager–Machlup functional weights paths; it does not, by itself, say which functional a given open system carries, in which variables, or with what metric. Those choices are exactly the modeler-supplied structure that Section 20 identified as the shared defect of all three traditions.

Fourth, the tradition never derived its own functional. At every stage the Lagrangian, the index, or the stochastic action is *written down* — from symmetry, from phenomenology, from the model at hand — and the variational principle then processes it. Nothing in the tradition says which functional a given system must carry; that question was not even well posed until there was an upstream object to project. The same gap will reappear, in different words, in both of the other traditions: MEPP cannot say which entropy production to extremize, and the free energy principle cannot say where the boundary of the believer lies.

Fifth, no classical result says why an open, driven, noisy system should follow any least-time path at all. Fermat’s principle is exact for rays; Hamilton’s for closed conservative systems; Feynman’s argument requires coherent interference; Onsager–Machlup requires the small-noise limit of a specific linear model. The trajectory gauge of this Part (Section 27) is precisely the missing derivation: it obtains a generalized least-time law for open systems as the small-noise reading of the projected functional, with the effective optical metric supplied by the projection rather than chosen by hand — and with the regime in which the reading holds stated as a theorem rather than discovered by failure.

Table 14: The action tradition’s own scope, row by row. Each row is one classical result, with the regime in which it is exact and the limitation that stops it at the boundary of open systems. The table’s point is the last column read downward: every exact result either excludes dissipation and noise outright or handles them as a probability weight whose functional must be supplied by hand — so the tradition, for all its power, contains no theorem that sends a driven, noisy, open system down a least-time path. That theorem is what the trajectory gauge of Section 27 supplies, and this table is the measure of what supplying it is worth.

Result	Exact where	Limitation
Fermat’s least time (1662)	Geometric optics; rays between fixed endpoints	Stationary, not always least; no matter, no noise
Maupertuis–Euler–Lagrange–Hamilton action (1744–1834)	Closed conservative mechanics; all of fundamental field theory	Excludes dissipation by construction; no open-system form (Bauer 1931)
Feynman’s stationary phase (1948)	Classical limit of coherent quantum evolution	Requires interference; decoherence and driving blur the extremum
Onsager–Machlup action (1953)	Small-noise limit of linear dissipative dynamics	A probability weight, not a dynamics; the functional and metric are inputs, not outputs

22 Minimum entropy production (Prigogine): history, applications, and limitations

22.1 A century of minimum-dissipation theorems

Minimum entropy production (MinEP) is often presented as a single theorem with a single author. It is in fact the last and sharpest member of a family of minimum-dissipation theorems running through a century of linear physics, and every member owes its minimum property to the same mathematical circumstance — a quadratic dissipation functional generated by a symmetric, positive-definite linear response. Kirchhoff proved in 1848 that the steady currents in a network of ohmic conductors minimize the total Joule heat among all distributions compatible with the sources (Kirchhoff, 1848); Helmholtz proved the analog for creeping viscous flow in 1868 (Helmholtz, 1868); Rayleigh introduced the dissipation function for damped vibration in 1873 (Rayleigh, 1873). In each case the minimum is a property of *linearity*, not of nature at large — the pattern that recurs through everything below.

Onsager’s contribution of 1931 must be stated carefully, because it is systematically confused with the theorem that came fourteen years later. Alongside the reciprocal relations $L_{ik} = L_{ki}$, Onsager formulated a principle of least dissipation: at *given forces*, the realized fluxes extremize $\sigma(J, X) - \Phi(J, J)$, where Φ is a quadratic dissipation function of the fluxes (Onsager, 1931). The stationarity condition of that functional *is* the linear constitutive law $J = LX$ — so Onsager’s principle is a flux-side encoding of the transport law itself, and says nothing about which *state* a driven system selects. Prigogine’s theorem, by contrast, varies the free forces at fixed clamped forces and characterizes the steady state. Keeping the two apart is not pedantry: their conflation generated much of the literature’s confusion, as Jaynes later documented (Jaynes, 1980), and the derivation of Section 28 recovers them as distinct readings of one stationary point.

The theorem itself arrived in 1945, in Prigogine’s communication to the Belgian academy from De Donder’s Brussels school, with the full exposition in his 1947 thesis (Prigogine, 1945, 1947): for linear phenomenological laws with constant, symmetric, positive-definite coefficients, and with a subset of the forces clamped by reservoirs, the stationary state is the state of least entropy production, and P decreases monotonically during the approach to it. The careful neutral statement of the hypotheses is Chapter V of de Groot and Mazur (1962). The first microscopic test came in 1954, when Klein and Meijer confirmed the theorem by master-equation methods for gas flow between two vessels — *to first order in the temperature difference* (Klein and Meijer, 1954). Already the verification carried the message, made precise half a century later, that MinEP is a first-order-in-driving statement.

22.2 The theorem, exactly

Theorem 22.1 (Prigogine, 1945). *Let forces $X_1, \dots, X_n$ and conjugate fluxes $J_1, \dots, J_n$ carry entropy production $P = \sum_i J_i X_i \geq 0$, and suppose (A1) linearity, $J_i = \sum_k L_{ik} X_k$; (A2) constancy of the L_{ik} ; (A3) Onsager reciprocity, $L_{ik} = L_{ki}$; (A4) positive definiteness of L ; and (A5) a clamped/free split: $X_1, \dots, X_m$ held fixed by reservoirs, the rest free. Then the state minimizing P over the free forces is exactly the stationary state, characterized by the vanishing of the fluxes conjugate to the free forces; and along the relaxation, $dP/dt \leq 0$, with equality only at the stationary state.*

The proof is a page: under (A1)–(A3), P is a positive-definite quadratic form in the forces, so $\partial P/\partial X_j = 2J_j$ for each free index j , and stationarity of P is the vanishing of the free fluxes — which is the definition of the steady state; positive definiteness makes the stationary point a minimum, and the same quadratic structure makes P a Lyapunov function of the relaxation. Every word of the theorem is load-bearing, and each hypothesis, when dropped, has a named counterexample waiting (Section 22.4).

22.3 The failed generalization, and what it created

The attempt to carry the theorem beyond the linear regime occupied Prigogine and Glansdorff for nearly two decades and ended in one of the more instructive negative results of twentieth-century physics. Splitting the time derivative of P into a force part and a flux part (Glansdorff and Prigogine, 1954), they proved in 1964 that beyond the linear regime only the force part retains a definite sign, $d_X P/dt \leq 0$ — the general evolution criterion — while the flux part, and hence dP/dt itself, does not (Glansdorff and Prigogine, 1964). The deeper disappointment, fully acknowledged by the authors, is that $d_X P$ is *not an exact differential*: outside the linear regime there exists, in general, no functional of the state whose descent the criterion expresses. Nothing is minimized. Far from equilibrium there is, provably, no potential of the entropy-production family (Glansdorff and Prigogine, 1971).

Out of the failure came the positive program. If far-from-equilibrium steady states obey no extremum, their stability must be studied directly — and it can be lost. Beyond a critical affinity the Brusselator’s homogeneous state destabilizes to Turing patterns or limit cycles (Prigogine and Lefever, 1968); past a critical Rayleigh number a heated fluid layer starts to convect; and the

selected regime is emphatically *not* the branch of least entropy production. These are the dissipative structures of the 1977 Nobel Prize, and the point is often garbled and worth stating plainly: the prize honors, in essence, Prigogine’s demonstration of the limits of his own 1945 theorem (Prigogine, 1978).

22.4 Critiques, counterexamples, and the exact scope

Three classical critiques fix the boundary.

Landauer’s blowtorch (1975). Heat a stretch of the transition path between two locally stable states, in a region where the stationary occupation is vanishingly small, and the relative occupation of the two states shifts — although no functional built from the entropy family registers the heating (Landauer, 1975a). The steady state is a property of the full kinetics, not of thermodynamic bookkeeping. In the companion paper the failure appears arbitrarily close to equilibrium: the fluctuating current through a driven inductor does not minimize entropy production (Landauer, 1975b). The modern diagnosis is time-reversal parity (Bruers et al., 2007): the theorem tacitly assumes all variables even under time reversal, and for odd variables — currents through inductors — the same linear regime yields a *maximum*-type characterization, so that neither slogan survives even in linear electrical networks.

Jaynes’s two-temperature network (1980). In a resistor network with the resistors held at different temperatures, the physical currents obey Kirchhoff’s laws and minimize the total *heat* production; a minimum *entropy* production principle would weight each resistor by $1/T_k$ and predict a different, wrong current distribution (Jaynes, 1980). Nonuniform temperature makes the quadratic form’s coefficients state-dependent across the system — hypothesis (A2) fails — and Jaynes’s verdict on scope stands: where the principle is exact it adds nothing beyond the linear laws that imply it, and where it would add something, it is wrong.

The stochastic-thermodynamics verdict. The master-equation setting makes the approximate character of MinEP a theorem rather than a grievance. With entropy production given by the Schnakenberg functional (Schnakenberg, 1976), Maes and Netočný proved that for dynamics obtained by a small deformation of detailed balance, the large-deviation rate functional of the occupation statistics is, to leading order, a quarter of the excess entropy production (Maes and Netočný, 2007) — so MinEP is *derived*, and simultaneously *bounded*, as the first order of a fluctuation functional. At the next order a time-symmetric contribution — the frenetic part, measuring traffic and dynamical activity — enters independently, and no functional of entropy production alone, minimized or maximized, determines the steady state (Maes and Netočný, 2013). Landauer’s blowtorch is the canonical illustration: the heating alters the frenetic factors while the entropic bookkeeping is blind to them.

One sentence carries this section into the derivation. Near equilibrium the theorem is exact, which means any deeper account of driven systems must reproduce it in the gentle-driving limit — just as any theory of motion must reproduce Newton’s at low speed — and the current gauge of Section 28 passes exactly this checkpoint, recovering Theorem 22.1, hypotheses and all, as the force-side reading of its stationary point.

Table 15: The exact scope of Prigogine’s theorem, one row per load-bearing hypothesis. Each row names the assumption in Theorem 22.1, the classical counterexample that lives at that door, and what the counterexample proves. The table should be read as a passport rather than an indictment: a theorem whose exceptions can all be mapped onto named hypotheses is delimited, not discredited — the healthiest condition a scientific statement can be in — and the derivation of Section 28 must, and does, reproduce both the theorem inside this passport and its silence outside it.

Hypothesis	Counterexample when dropped	What it proves
(A1) Linear laws	Brusselator beyond critical affinity; Bénard convection	No entropy-production potential exists far from equilibrium; selection passes to stability theory
(A2) Constant coefficients	Jaynes’s two-temperature network	The physical state minimizes heat, not entropy production; MinEP predicts the wrong currents
(A3)–(A4) Symmetric positive response Evenness under time reversal (tacit)	Loss of reciprocity (magnetic fields, active media) Landauer’s driven inductor	The quadratic structure that carries the proof is absent Odd variables convert the minimum into a maximum-type principle within the same linear regime
(A5) Clamped/free split	Nothing clamped	Nothing is selected; the theorem has no content

23 Maximum entropy production: history, applications, and limitations

23.1 From least dissipation to a maximum principle

Every modern statement of maximum entropy production descends, ironically, from the minimum principles of the last section: Ziegler’s formulation is the Legendre-dual face of the same quadratic structure that carries Kirchhoff, Rayleigh, and Onsager. Working in continuum thermomechanics and plasticity between 1958 and 1983, Ziegler posited the *orthogonality principle*: the dissipative force realized by a material is normal to the level surface of a prescribed dissipation function of the fluxes — equivalently, at prescribed forces the fluxes maximize the entropy production subject to the single constraint that production equal the power expended (Ziegler, 1961, 1983). In rate-independent plasticity this is the classical normality rule — plastic strain rates normal to the yield surface — where it has done quiet, honest service for decades. For quadratic dissipation the orthogonality condition returns exactly Onsager’s linear law, so in the linear regime Ziegler’s maximum and Onsager’s least dissipation are one statement in two conventions; and Polettini has shown that for non-homogeneous dissipation functions the orthogonality condition and the constrained maximization come apart, so that “Ziegler’s principle” is a convention-dependent family rather than one principle (Polettini, 2013).

One member of the family is a theorem in the full sense, and it is worth naming exactly because it is so often cited on the principle’s behalf: Kohler (1948) and Ziman showed that the stationary solution of the *linearized* Boltzmann equation maximizes the entropy production among

trial distributions satisfying a self-consistency constraint (Ziman, 1956). The proof is three lines of Cauchy–Schwarz in the collision operator’s inner product — and the theorem lives entirely in linearized transport, where it restates least dissipation and adds nothing beyond the transport law it encodes.

23.2 Paltridge’s climates: the empirical core

What made MEPP famous is none of the above; it is a 1975 discovery in climate science. Paltridge partitioned the Earth into zonal boxes exchanging heat, wrote only the steady energy balances — shortwave in, longwave out, an unknown meridional transport — and closed the underdetermined system by demanding that the material entropy production of the transport be maximal (Paltridge, 1975, 1978). In the two-box form the whole construction fits on a postcard, and it is worth having on the page, because it is the strongest empirical card the principle holds. Let box 1 (low latitudes) absorb shortwave power I_1 and box 2 absorb $I_2 < I_1$; let $F \geq 0$ be the heat carried poleward; and linearize the outgoing longwave emission as $A + BT$. The steady balances

$$I_1 - (A + BT_1) - F = 0, \quad I_2 - (A + BT_2) + F = 0 \quad (23.1)$$

make each temperature affine in the single unknown F : the sum fixes the mean temperature $\bar{T}$ independently of the transport, and the difference gives the contrast $\Delta T(F) = (\Delta I - 2F)/B$ with $\Delta I \equiv I_1 - I_2$. The entropy production of the transfer,

$$\sigma(F) = F \left(\frac{1}{T_2} - \frac{1}{T_1} \right) = \frac{F \Delta T(F)}{T_1(F) T_2(F)}, \quad (23.2)$$

vanishes at both ends of the admissible range — at $F = 0$ (finite force, no flux) and at $F_{\text{eq}} = \Delta I/2$, where the transport has erased the very temperature difference that drives it (finite flux, no force) — so it has an interior maximum. In the small-contrast regime $\Delta T \ll \bar{T}$ the curve is an exact parabola and the maximum sits at

$$F^* = \frac{\Delta I}{4}, \quad \Delta T^* = \frac{\Delta I}{2B} = \frac{\Delta T_{\text{rad}}}{2}, \quad \sigma_{\text{max}} = \frac{(\Delta I)^2}{8B\bar{T}^2} : \quad (23.3)$$

the selected state transports exactly half the equalizing flux and sustains exactly half the radiative-equilibrium contrast. The trade-off in (23.2) — every unit of flux erodes the force that pays for it — is the entire mathematical content of the climate MEPP in its simplest form, and the principle bets that nature sits at the optimum. Fed nothing but absorbed sunlight and the two empirical radiation constants, the closure reproduces the observed equator-to-pole contrast — and when Lorenz applied the identical construction to Mars and Titan it worked there too, with no adjustable transport coefficient, while a diffusive closure tuned to Earth fails when exported (Lorenz et al., 2001, Ozawa et al., 2003). Three planets, one postcard formula: this is the empirical core of the MEPP literature, and any assessment of the principle has to reckon with it.

It must equally reckon with what the closure ignores. It contains no rotation rate, no momentum budget, no moist processes — Goody pressed exactly this against the theoretical program (Goody,

2007); the maximized quantity is the material entropy production while the planet’s dominant, radiative production is not extremized (Essex, 1984); and the prediction depends on how the system is partitioned and which fluxes are treated as free — the resolution and constraint-choice dependence emphasized in the modern consensus assessment (Singh and O’Neill, 2022). None of this erases the result; it sharpens what the result is: a parameter-free regularity that holds across three atmospheres is a fact about atmospheres, however innocent the formula — a benchmark that any candidate theory, including this Part’s, must reproduce or explain away.

23.3 The claimed derivation, and where it broke

A general proof of MEPP was claimed in 2003–2005. Dewar applied Jaynes’s maximum-entropy inference to trajectories: assign path probabilities $p_\Gamma = Z^{-1}e^{A_\Gamma}$ by maximizing path entropy subject to prescribed mean boundary fluxes and local conservation laws (Dewar, 2003, 2005). The construction’s early steps are sound and even fruitful. Integrating the constraints by parts splits the action into a part even under time reversal (depending only on the endpoints) plus one half of the time-integrated entropy production,

$$A_\Gamma = A_\Gamma^{\text{even}} + \frac{\tau \sigma_\Gamma}{2k_B}, \quad (23.4)$$

from which the transient fluctuation theorem $p_\Gamma/p_{\Gamma_R} = e^{\tau\sigma_\Gamma/k_B}$ — and with it the second law, as a Jensen corollary — follows immediately, since reversing a path flips the odd half and preserves the even one. That much is uncontested. The failure is localized in the closing steps, and the critical literature located it precisely. Dewar’s 2003 route requires the number of paths at given irreversible action to be an *increasing* function of that action — an assumption nowhere derived; Bruers exhibited an explicit lattice-current model in which the count runs the other way, so that the same machinery, run honestly in the linear regime, delivers Prigogine’s *minimum* (Bruers, 2007). The 2005 route rests on a Gaussian saddle-point evaluation that is valid only for weak flux fluctuations — that is, near equilibrium — as Grinstein and Linsker showed, which forecloses the claimed far-from-equilibrium generality (Grinstein and Linsker, 2007). Dewar conceded both points in print and reframed the principle as an inference algorithm whose output is only as reliable as the constraints fed into it — no longer a claimed theorem of physics (Dewar, 2009).

Bruers’s taxonomy is worth carrying forward: at least five inequivalent statements circulate under the one name — the linear-response theorem (Kohler–Ziman, equivalent to least dissipation), Paltridge’s partial steady-state closure, a total steady-state maximization, non-variational selection among discrete far-from-equilibrium states, and the near-equilibrium minimum that Dewar’s own machinery actually implies (Bruers, 2007). The climate successes are not underwritten by the path-information argument; they stand, underived, on their own.

23.4 Scope, and the distinction that prevents most confusion

The landscape in one sentence: MEPP is one exact theorem (linearized transport, where it adds nothing), one durable geophysical heuristic (Paltridge’s, sharpened by three planets, still underived), and one failed general derivation (broken at a named step and conceded) — while the rigorous

far-from-equilibrium variational structure that actually exists, in macroscopic fluctuation theory, is a corrected *minimum*-dissipation principle for a functional that is not entropy production alone (Bertini et al., 2015).

A final distinction prevents most of the confusion the slogan invites. That every real process produces entropy is *bookkeeping* — computed from the flows the way a balance is computed from transactions, automatically true. MEPP claims something categorically stronger: that among all flow patterns the constraints allow, nature *selects* the one of largest production. The ledger is true by construction; the selection is a physical hypothesis, and it is the hypothesis, not the ledger, that has no general proof. The current gauge of Section 28 derives the selection — as the flux-side reading of one stationary point, valid exactly where the curl and gauge sectors do no ledger work, with the force-side reading of the same point recovering Prigogine’s minimum — and the failure theorem of Section 30 marks where both readings stop.

Table 16: The status of maximum entropy production, separated into the three things that circulate under its name. Each row gives the statement, its strongest verified support, and its present standing in the technical literature. The separation is the section’s whole point: the theorem is confined to linearized transport where it restates least dissipation; the planetary heuristic is empirically striking and underived, its predictions dependent on the modeler’s choice of constraints and functional; and the claimed general derivation failed at two named steps and was conceded by its author. A reader who keeps the three rows apart will not be misled by literature that cites the first and third on behalf of the second.

Statement	Strongest support	Standing
Linear-response maximum (Kohler–Ziman; Ziegler, quadratic case)	Three-line proof for the linearized Boltzmann equation (Ziman, 1956)	Theorem; equivalent to least dissipation; adds nothing beyond the linear law
Paltridge’s climate closure	Equator-to-pole contrast on Earth, Mars, Titan with no tuned coefficient (Lorenz et al., 2001)	Underived heuristic; constraint- and resolution-dependent; ignores the dominant radiative budget
General far-from-equilibrium principle (Dewar 2003–2005)	Path-information construction through the fluctuation theorem	Refuted at the monotonicity and Gaussian steps (Bruers, 2007, Grinstein and Linsker, 2007); conceded (Dewar, 2009)

24 The free energy principle: history, applications, and limitations

24.1 Three streams, then a fourth

The free energy principle did not appear from nowhere in 2005; it is the confluence of three older streams, with a fourth grafted on late — and because the derivational disputes below turn on *which* stream a given claim belongs to, the history matters.

The first stream is perceptual epistemology. Helmholtz’s doctrine of unconscious inference — percepts are conclusions drawn from sensory premises together with prior experience — is the conceptual ancestor claimed by the whole tradition (Helmholtz, 1867); the connection is programmatic, since Helmholtz supplies no objective functional. The second is cybernetic regulation:

Ashby framed adaptive behavior as the maintenance of essential variables within bounds (Ashby, 1952), and the good-regulator theorem of Conant and Ashby — every good regulator of a system must be a model of that system — is the FEP’s central claim restated in advance (Conant and Ashby, 1970). The third is variational statistics. The quantity later called F entered neural modeling through machine learning: the Helmholtz machine trains a recognition density to bound the log evidence of a generative model (Dayan et al., 1995), and Neal and Hinton showed that expectation–maximization is coordinate descent on a single free energy, so that inference and learning optimize one objective (Neal and Hinton, 1998). Predictive coding gave the tradition its canonical cortical algorithm — descending predictions, ascending prediction errors (Rao and Ballard, 1999) — and it is worth saying plainly that predictive coding neither requires nor implies the FEP, a distinction that matters when empirical support for the former is offered on behalf of the latter.

Friston assembled the streams into a principle between 2005 and 2010: hierarchical cortical processing as empirical-Bayes inference minimizing free energy (Friston, 2005), action added as a second route to the minimum — change the predictions to match the sensations, or change the sensations to match the predictions (Friston et al., 2006) — and the canonical statement in 2010 (Friston, 2010). Active inference followed (Friston et al., 2010, 2017), consolidated in textbook form (Parr et al., 2022); this stratum now runs a substantial fraction of computational neuroscience and computational psychiatry.

The fourth stream is the *physics phase*. From 2013 Friston pivoted from brains to matter: a random dynamical system possessing a Markov blanket at nonequilibrium steady state will “appear to” perform Bayesian inference about what lies beyond its blanket (Friston, 2013), a program culminating in the 2019 monograph (Friston, 2019). This is the stratum that claims a law of nature — and it is the stratum on which the critiques fell.

24.2 The mathematics that is solid

The core is textbook variational inference, uncontested by any critic. With a generative density $p(s, \psi)$ over sensory data s and latent states ψ , and a recognition density $q_\mu(\psi)$ parameterized by internal states μ , the variational free energy $F(s, \mu) = \mathbb{E}_{q_\mu}[\ln q_\mu(\psi) - \ln p(s, \psi)]$ admits three decompositions, each carrying one reading of the principle:

$$F(s, \mu) = D_{\text{KL}}[q_\mu(\psi) \parallel p(\psi \mid s)] - \ln p(s) \quad (24.1)$$

$$= \underbrace{\mathbb{E}_{q_\mu}[-\ln p(s, \psi)]}_{\text{energy}} - \underbrace{H[q_\mu]}_{\text{entropy}} \quad (24.2)$$

$$= \underbrace{D_{\text{KL}}[q_\mu(\psi) \parallel p(\psi)]}_{\text{complexity}} - \underbrace{\mathbb{E}_{q_\mu}[\ln p(s \mid \psi)]}_{\text{accuracy}}. \quad (24.3)$$

Since the divergence in (24.1) is nonnegative, $F \geq -\ln p(s)$, with equality exactly when q_μ is the posterior: minimizing F over μ makes the internal states approximate Bayesian inference (perception); minimizing F over actions that change s is the additional move that defines active inference; and $-F$ is the evidence lower bound of machine learning, descended through Dayan et al. (1995) and Neal and Hinton (1998) from the variational bounds of statistical mechanics. In

hierarchical Gaussian models, gradient descent on F is precision-weighted prediction-error dynamics,

$$\dot{\mu} \propto -\frac{\partial F}{\partial \mu} = (\partial_{\mu} g)^{\top} \Pi \varepsilon, \quad \varepsilon \equiv s - g(\mu), \quad (24.4)$$

descending predictions $g(\mu)$ against ascending errors ε weighted by precisions Π — the predictive-coding process model — with attention entering as the optimization of the precisions themselves, since stationarity of F in Π sets each inverse precision to the averaged outer product of its errors. All of this is standard, empirically fruitful against neural data, and stands independently of any claim about matter in general.

24.3 The contested stratum, and the critiques that located it

The physics phase asserts a chain: a stationary system with a sparse-coupling boundary (the Markov blanket) factorizes its stationary density; the blanket induces a synchronization map from internal to external averages; and the expected internal states then flow down the gradient of a free energy — so that persistence itself entails inference. Each link has been examined, and the examination is the sharpest technical literature in this Part.

Biehl, Pollock, and Kanai proved that the two blanket definitions in circulation — sparsity of the flow’s Jacobian, and factorization of the stationary density — are logically independent; that the rewriting of the expected internal flow requires an absence of solenoidal coupling across the partition that is nowhere stated; and that the monograph’s Free Energy Lemma fails at face value, with counterexamples (Biehl et al., 2021). The authorized reply conceded the observations and supplied a sparse-coupling ansatz — sufficient conditions under which the argument goes through — thereby restricting the principle’s antecedents rather than rescuing its generality (Friston et al., 2021). Aguilera, Millidge, Tschantz, and Buckley then posed the decisive quantitative question: on what fraction of parameter space do the conditions hold, in the linear (Ornstein–Uhlenbeck) model class the FEP literature itself uses? The answer: a narrow subspace, whose defining symmetry amounts to the *absence of perception–action asymmetry* — the very asymmetry living systems exhibit — and even where the construction completes, the inferential reading adds no predictive content beyond the dynamics themselves (Aguilera et al., 2022). Complementary critiques established that in applications the blanket is induced by the modeler’s parcellation rather than discovered in the system (Raja et al., 2021); that the move from Pearl’s epistemic blanket to a physical boundary of agents is a category transition the mathematics does not underwrite (Bruineberg et al., 2022); and that a principle defended as unfalsifiable-by-status — only its applications can fail — cannot function as an empirical law of biology (Colombo and Wright, 2021).

The program’s response has been a retreat executed with candor. The rigorous results that exist are confined to linear-Gaussian settings (Da Costa et al., 2021); the 2023 reformulation abandons the stationary density factorization for a path-integral statement in which the principle holds by construction for systems admitting a particular partition (Friston et al., 2023a); and the programmatic statement of the resulting Bayesian mechanics presents the FEP as a principle of least action *on belief spaces* — not itself falsifiable, but a scheme for generating falsifiable process models (Ramstead et al., 2023). The instrumentalist reading secures the framework by surrendering

the claim that made it famous. One cannot simultaneously advertise a theory of every persisting thing and retreat, under technical pressure, to the position that the formalism is a lens the modeler may choose to wear.

24.4 The honest summary, and what a derivation owes it

Three assumptions beneath the Gaussian machinery are worth naming even inside the solid stratum, because each has a characteristic biological failure mode: a posterior summarized by a single mode (perceptual bistability defeats it), an adequate second-order expansion (strong nonlinearity in g defeats it), and quadratic error penalties (heavy-tailed sensory noise strains them). These are scope conditions of the *algorithm*, cataloged within the tradition itself; the inferential gauge below inherits none of them as postulates, because its metric and its boundary arrive from the projection rather than from the model class.

For a physicist, the assessment of the mid-2020s compresses to one sentence: the bound is trivial, the decomposition is classical, and the bridge from blankets to inference is either false in general, true by construction, or true under fine-tuned conditions — depending on which paper one reads. What the tradition lacks is exactly what it cannot supply from inside: a boundary that is *derived* rather than drawn by the modeler, a metric on beliefs that is *forced* rather than chosen, and a regime theorem that says where the inferential reading holds. The inferential gauge of Section 29 supplies all three — the blanket from the declared ports of the packet, the Fisher metric from the ledger’s own informational term, and the rank and sector conditions as theorems — and Section 30 states where the reading fails, which no formulation of the FEP has ever done for itself.

Table 17: The free energy principle separated into its four strata, with the verdict the technical literature has reached on each. The first two rows are uncontested and empirically productive; every controversy lives in the third and fourth. The table enforces the distinction that most discussion of the principle blurs: a model that fits brains is judged as a model, by prediction and failure against data, while a law about matter must hold for everything that persists — and the two kinds of claim, frequently interleaved in the same paper, have opposite standings in the record. The derivation of Section 29 engages the fourth stratum on its own terms and replaces its assumptions with derived structure.

Stratum	Content	Verdict
Variational bound	$F \geq -\ln p(s)$; ELBO (24.1)	Textbook; uncontested (Dayan et al., 1995, Neal and Hinton, 1998)
Process models	Predictive coding; active inference; precision as attention	Legitimate, practical, empirically productive (Friston et al., 2017, Rao and Ballard, 1999)
Physics phase	Blanket $\Rightarrow$ factorization $\Rightarrow$ synchronization $\Rightarrow$ gradient flow	Non-generic; silent assumptions; counterexamples (Aguilera et al., 2022, Biehl et al., 2021)
Bayesian-mechanics retreat	Least action on belief spaces; holds by construction	Instrumentalist; not falsifiable as stated (Colombo and Wright, 2021, Friston et al., 2023a, Ramstead et al., 2023)

25 Constitutional statement of the Qi-EP principle

The conceptual claim of Qi-EP should be stated before any detailed derivation:

Definition 25.1 (Qi-EP). For any admissible open subsystem S embedded in the GG-6 bulk and observed through OGGEF, there exists a projected extremal functional

$$\Omega_{\text{Qi-EP}}^{(S)} = \Phi_{\text{OGGEF}}^{(S)}[S_{\text{GG-6}}] \Big|_{\Gamma_{\text{adm}}} \quad (25.1)$$

such that the realized boundary dynamics of S satisfy

$$\delta\Omega_{\text{Qi-EP}}^{(S)} = 0 \quad \text{on} \quad \Gamma_{\text{adm}}. \quad (25.2)$$

Under three natural boundary gauges this same stationary condition appears respectively as PLT, MEPP, and FEP.

Three comments are essential.

First, the word “projected” is not decorative. The boundary principle is not postulated independently. It is induced from the 6D GG-6 action through the open-system projection map already established in GG-A5-GDOS (Arisaka, 2026, A2-GG6; Arisaka, 2026, A5-GDOS). This means that the extremal selector on the boundary is constitutionally tethered to the same ultraviolet structure that fixed BI-6, the OGN integers, and the current ratios.

Second, the phrase “under three natural boundary gauges” is also literal. The projected functional may be described in terms of trajectories, in terms of currents and forces, or in terms of inferential states on a statistical manifold. These are not three unrelated principles living side by side. They are three descriptions of the same projected geometry under three different coarse-grained coordinate choices.

Third, Qi-EP is not restricted to quantum mechanics, nor to biology, nor to any single domain. The same projected functional can describe a geodesic-like transport path in one application, a non-equilibrium stationary flux in another, and an inferential or adaptive update in a third. This is exactly what a universal open-system selector must do if it is to serve as Pillar II of a general theory.

Dual-layer architecture and the compatibility relation $\mathfrak{C}_{\text{EP}}$. Qi-EP is, in the L16 language, a dual-layer principal pillar. Its *algebraic layer* (T-EP.1) is the unified gauge theorem (Theorem 32.1): the same projected boundary functional $F_{\text{GFF}} = F_{\text{GEP}} + F_{\text{CEP}}$ governs three apparently-distinct variational principles simultaneously, with no contradiction. Its *operational/physical layer* (T-EP.2) is the three named gauges PLT (Theorem 27.1, the trajectory gauge of Maupertuis–Hamilton), MEPP (Theorem 28.1, the current gauge of Onsager–Prigogine–Paltridge), and FEP (Theorem 29.1, the inferential gauge of Friston), each providing a different falsifiable laboratory prediction. The two layers are tied together by the named compatibility relation

$$\boxed{\mathfrak{C}_{\text{EP}} \equiv \text{T-EP.1} \implies \text{T-EP.2} \quad (\text{via the boundary projection of the 6D bulk action functional}).} \quad (25.3)$$

The unified theorem at the algebraic level forces the three gauges to share one projected functional; each gauge then provides a different falsifier on a different physical regime (trajectories at high visibility, currents at near-equilibrium, inference at strong noise).




Figure 8: How three principles that look like rivals turn out to be one. The single projected functional at the top is the boundary image of the six-dimensional least-action principle; it is not a new postulate. Each arrow is the same stationarity condition asked with a different quantity held primary, and the label on the arrow names which: hold the path and the condition returns a least-time law, hold the currents and it returns the flux-force relation of irreversible transport, hold the beliefs and it returns the natural-gradient update used in the brain sciences. The bottom row names the community that has been using each form. The dashed link records what the unification settles on the way through: maximum entropy production and Prigogine’s minimum entropy production are conjugate readings of one stationary point under dual control conventions — the first fixing the transport law, the second the state — not competing claims about nature. One system may occupy all three gauges at once — a cell following a gradient traces a path, maintains fluxes, and corrects an estimate.

25.1 The contract: five criteria a genuine derivation must meet

The histories of Sections 21–24 fix the specification, and it is stated here as a contract so that the reader can hold the rest of the Part to it. A derivation worth the name must: (i) produce the three principles from *one* object that is itself derived — the projection of an upstream action, not a fresh postulate; (ii) deliver the restrictions as *theorems* — the linear-regime confinement of MinEP, the parity and inhomogeneity failures, the coupling and rank conditions of the inferential reading — rather than discovering them by failure; (iii) dissolve the minimum–maximum quarrel, in agreement with the resolution the critical literature reached from below; (iv) replace every modeler-chosen structure with derived structure — the choice of *which* entropy production by the ledger fixed upstream, the blanket by the declared ports of the packet, the synchronization map by a rank condition; and (v) state the conditional and open elements explicitly, and nowhere waive them. The sections that follow discharge the contract clause by clause: (i) in Section 26, (ii) in Sections 27–29 and 30, (iii) in Section 28, (iv) in Sections 26 and 29, and (v) in Section 32. A framework that “derived” MEPP as universally true would be wrong — the counterexamples are theorems too. What

a correct derivation must deliver is the functional whose restricted readings the three principles are, together with the location of the restrictions.

26 From the 6D bulk action to the projected boundary action

26.1 The imports, with their grades

Three objects are taken from upstream, and nothing else is assumed; the derivation’s honesty is easiest to audit if they are listed before they are used. First, the *boundary law*: the GDOS continuity equation with its four-sector current — gradient, curl, noise, and gauge, Eq. (4.4) — derived in GG-A5-GDOS from the projection of the bulk through the open-system channel (Arisaka, 2026, A5-GDOS), and imported here at grade Derived. Second, the *ledger*: the single scalar functional $F_{\text{GFF}} = F_{\text{GEP}} + F_{\text{CEP}}$ with its balance identity, established in GG-A6-GRIP (Arisaka, 2026, A6-GRIP) and imported at grade Derived; every appearance of “the entropy production” below is an appearance of this ledger’s dissipation term, which is what removes the choice of functional that the MEPP literature never managed to fix (Section 23). Third, the *packet*: the SIP declaration of Part II, whose ports are constitutive data — declared when the system is written down, not discovered inside it — and which will do for the inferential gauge what walls do for thermodynamics. The bulk action itself and the OGGEp projection that carries it to the boundary enter from GG-A2-GG6 and GG-A1-Constitution at the grades their home papers assign. Everything else in this Part is constructed on the page.

26.2 The mathematical tools, briefly

A reader arriving from the life sciences needs four instruments, each stated here in one breath and used repeatedly below. *Path weights*: for noisy dynamics the object that replaces the trajectory is a probability over trajectories, and in the weak-noise regime that probability concentrates as $P[z] \asymp e^{-\mathcal{A}[z]/\varepsilon}$ — so extremizing an action and finding the most probable history become the same act (Onsager and Machlup, 1953, Touchette, 2009). *Large deviations*: the same concentration statement for densities and currents, which is what lets one functional govern a single swimming cell and a whole population’s flux pattern. *Legendre duality*: a convex stationary point can be described with either of two conjugate variables held fixed, and the two descriptions extremize different functions of the same point — the entire minimum–maximum quarrel of Section 23 lives inside this one fact. *Fisher geometry*: a family of probability densities is a curved manifold whose natural metric is the Fisher information, and gradient descent in that metric — natural-gradient descent — is the coordinate-free way to go downhill on beliefs (Amari, 1998). None of these is new mathematics; what is new in this Part is where the functional, the convex form, and the metric come *from*.

26.3 The projection, step by step

We begin from the ultraviolet starting point already established in GG-A2-GG6: the six-dimensional bulk with total action

$$S_{\text{GG-6}}[\Psi] = \int_{\mathcal{M}_6} d^6x \sqrt{-g_6} \mathcal{L}_6(\Psi, \nabla\Psi), \quad (26.1)$$

where Ψ collectively denotes the 6D metric, gauge, Green–Schwarz, and matter fields. By construction the bulk dynamics satisfy

$$\delta S_{\text{GG-6}} = 0, \quad (26.2)$$

together with the BI-6 consistency constraints. Equation (26.2) is the constitutional PLA of the GG-6 stage (Arisaka, 2026, A2-GG6, Feynman, 1948, Hamilton, 1834).

To derive the boundary selector, split the full field content into observer-accessible slow variables z and hidden or traced-over variables y :

$$\Psi = (z, y). \quad (26.3)$$

The symbol z should be interpreted broadly. Depending on the later application it may denote a trajectory coordinate, a flux coordinate, a reaction-network coordinate, a neural or cognitive state coordinate, or a collective macroscopic descriptor. The hidden variables y include KK modes, ultraviolet environmental modes, unresolved microstates, and any degrees of freedom eliminated by the boundary observer.

Concretely, the standard way to write “integrate out hidden degrees of freedom while keeping a variational structure” is the Feynman–Vernon influence functional, which we now recall as Eq. (26.4). Readers from biology or chemistry may think of this as the path-integral parallel of the partial-trace operation of Eq. (1.1): both objects discard the environment and keep only the structure relevant to the observer, but the path-integral form is more convenient when one wants to expose the variational principle hidden inside the dynamics.

The standard influence-functional representation of the reduced dynamics is then

$$\exp\left(\frac{i}{\hbar} S_{\text{IF}}[z, z']\right) = \int \mathcal{D}y \mathcal{D}y' \exp\left[\frac{i}{\hbar} \left(S_{\text{GG-6}}[z, y] - S_{\text{GG-6}}[z', y']\right)\right] \rho_y(0), \quad (26.4)$$

where, in plain words, S_{IF} is the effective action for the system variables z obtained by integrating over the environment variables y ; the two-sided integration over y and y' reflects the two-index structure of the density matrix (the closed-time-path doubling familiar from Schwinger–Keldysh formalism), and $\rho_y(0)$ is the initial state of the environment. The equation is the path-integral expression of the same reduced evolution usually described by CPTP maps, Nakajima–Zwanzig kernels, or GKLS generators (Breuer and Petruccione, 2002, Feynman and Vernon, 1963, Gorini et al., 1976, Lindblad, 1976). In the present paper, however, the decisive point is not merely that reduction exists; it is that the reduction inherits a variational structure from the bulk.

In the semiclassical, weak-noise, or large-deviation regime one may write

$$P[z] \asymp \exp \left[-\frac{1}{\varepsilon_{\text{eff}}} \mathcal{A}_{\text{eff}}[z] \right], \quad (26.5)$$

where ε_{eff} is the effective coarse-graining scale and $\mathcal{A}_{\text{eff}}$ is the reduced action functional (Onsager and Machlup, 1953, Touchette, 2009, van Kampen, 2007). More generally, when both densities and currents are retained explicitly, the most useful projected object is the current-density functional

$$\mathcal{A}_{\text{eff}}[\rho, J] = \int_0^T dt \left[\frac{1}{4} \int_{\mathcal{X}} (J - \rho u)^a (\rho D)_{ab}^{-1} (J - \rho u)^b dz + \mathcal{V}[\rho] \right], \quad (26.6)$$

subject to the continuity equation

$$\partial_t \rho + \nabla_a J^a = 0. \quad (26.7)$$

Here $u^a(z, t)$ is the projected drift field, $D^{ab}(z, t)$ is the symmetric positive diffusion/metric tensor inherited from the projected bulk geometry, and $\mathcal{V}[\rho]$ is the projected potential or constraint functional. Equation (26.6) is the single most important formula of the present section. It is the master boundary action from which the three gauges will be obtained.

Proposition 26.1 (Projected stationary-action proposition). *Within the GG-6/GDOS framework, the open-system boundary histories selected by OGGEF are stationary histories of the projected action $\mathcal{A}_{\text{eff}}[\rho, J]$ on the admissible class Γ_{adm} .*

Proof. Fix (ρ, J) and consider the hidden-variable extremal $y = \bar{y}[z]$ obtained by solving the hidden Euler–Lagrange equations at fixed boundary history. At tree level or in the leading large-deviation approximation one has

$$\mathcal{A}_{\text{eff}}[z] = \Re S_{\text{GG-6}}[z, \bar{y}[z]] + (\text{coarse-graining counterterms}). \quad (26.8)$$

Varying with respect to z gives

$$\frac{\delta \mathcal{A}_{\text{eff}}}{\delta z} = \left. \frac{\delta S_{\text{GG-6}}}{\delta z} \right|_{y=\bar{y}[z]} + \left. \frac{\delta S_{\text{GG-6}}}{\delta y} \right|_{y=\bar{y}[z]} \frac{\delta \bar{y}}{\delta z} + \delta(\text{counterterms}). \quad (26.9)$$

The second term vanishes because $\bar{y}[z]$ is by construction a hidden extremal. The remaining terms combine into the reduced stationarity condition. The current-density form (26.6) is the corresponding open-system representation in which the hidden modes appear only through the projected drift, diffusion, and potential data. Therefore the realized boundary history is a stationary history of $\mathcal{A}_{\text{eff}}$. $\blacksquare$

The proposition is conceptually decisive. It says that Qi-EP is not a separate empirical principle pasted onto the boundary. It is the projected variational shadow of the same 6D least-action law that governed the ultraviolet bulk. The remaining task is to show how the familiar boundary forms PLT, MEPP, and FEP arise from three different reductions of Eq. (26.6).

27 Trajectory gauge: derivation of PLT

The first gauge is the trajectory gauge. In this description the boundary packet remains sharply localized around a representative path $z(t)$, so that

$$\rho(\zeta, t) \approx \delta(\zeta - z(t)), \quad J(\zeta, t) \approx \delta(\zeta - z(t)) \dot{z}(t). \quad (27.1)$$

Substituting Eq. (27.1) into Eq. (26.6) gives the reduced trajectory action

$$\mathcal{A}_{\text{trj}}[z] = \int_0^T dt \left[\frac{1}{4} (\dot{z} - u(z, t))^a D_{ab}^{-1}(z, t) (\dot{z} - u(z, t))^b + V(z, t) \right], \quad (27.2)$$

where $V(z, t)$ is the local representative of $\mathcal{V}[\rho]$ on the sharply localized slice.

It is convenient to introduce the effective metric

$$g_{ab}(z, t) := \frac{1}{2} D_{ab}^{-1}(z, t) \quad (27.3)$$

and the associated one-form

$$a_a(z, t) := g_{ab}(z, t) u^b(z, t). \quad (27.4)$$

Then Eq. (27.2) becomes

$$\mathcal{A}_{\text{trj}}[z] = \int_0^T dt \left[\frac{1}{2} g_{ab} \dot{z}^a \dot{z}^b - a_a \dot{z}^a + \frac{1}{2} g_{ab} u^a u^b + V \right]. \quad (27.5)$$

Up to the effective scalar potential

$$U_{\text{eff}}(z, t) := \frac{1}{2} g_{ab} u^a u^b + V(z, t), \quad (27.6)$$

this is the standard quadratic mechanical form on a curved configuration manifold.

The Euler–Lagrange equations are

$$\frac{d}{dt} (g_{ab} \dot{z}^b - a_a) - \frac{1}{2} \partial_a g_{bc} \dot{z}^b \dot{z}^c + \partial_a a_b \dot{z}^b - \partial_a U_{\text{eff}} = 0. \quad (27.7)$$

Raising the free index and rewriting with the Levi–Civita connection of g_{ab} gives

$$\frac{D_g \dot{z}^a}{dt} = g^{ab} (F_{bc} \dot{z}^c - \partial_b U_{\text{eff}}), \quad F_{ab} := \partial_a a_b - \partial_b a_a. \quad (27.8)$$

Equation (27.8) has the expected interpretation: the projected path follows a geodesic-like motion in the effective boundary metric, corrected by a Lorentz-like curvature term and a projected scalar force.

To make contact with the language of PLT, fix an energy shell E and perform the Jacobi reduction. One obtains the generalized optical-time functional

$$\mathcal{T}[z] = \int_\gamma \left[\sqrt{2(E - U_{\text{eff}}(z))} g_{ab}(z) dz^a dz^b - a_a(z) dz^a \right]. \quad (27.9)$$

The stationary condition $\delta\mathcal{T} = 0$ is the boundary PLT form of the same projected action. The phrase “least time” should therefore be interpreted in the generalized sense appropriate to open systems: not a bare Euclidean travel time, but the optical or information-geometric time induced by the projected metric, the drift connection, and the effective potential.

Theorem 27.1 (PLT as the trajectory gauge of Qi-EP(T-Qi4-10.3)). *On the sharply localized trajectory slice $\mathcal{S}_{\text{trj}} \subset \Gamma_{\text{adm}}$, the projected current-density functional (26.6) reduces to the generalized optical-time functional (27.9). Its stationary histories satisfy Eq. (27.8). Hence PLT is the trajectory gauge of Qi-EP.*

Proof. The reduction from Eq. (26.6) to Eq. (27.2) follows directly by substitution of Eq. (27.1). Completing the square and defining Eqs. (27.3)–(27.6) gives Eq. (27.5). The Euler–Lagrange variation yields Eq. (27.7), which is equivalent to Eq. (27.8). Restricting to a fixed energy shell and performing the standard Jacobi reduction turns the quadratic action into the time functional (27.9). Therefore the stationary condition of the projected action is precisely the PLT condition in this gauge. ■

The physical interpretation is straightforward. The 6D bulk obeys PLA. Once the inaccessible directions and hidden environmental couplings are projected away, the observer sees a boundary packet moving in an effective metric with an effective connection and an effective scalar landscape. The realized path is then the one of stationary generalized travel time. In transport problems, search problems, signaling problems, and morphogenetic path-selection problems, this is the most natural boundary language.

Table 18: Step-by-step structure of the least-time derivation. Six steps run down the first column, from the 6D constitutional law to the operational least-time form, with the equation and its meaning alongside. This is the first of three parallel chains — the other two are Tables 21 and 23 — and the three are laid out identically on purpose, so that a reader can see that they differ only in which variable is held as primary.

Step	Equation	Meaning
6D constitutional law	(26.2)	The bulk obeys $\delta S_{\text{GG-6}} = 0$.
Projected boundary action	(26.6)	Hidden modes induce drift, diffusion, and a projected potential.
Trajectory slice	(27.1)	The packet is sharply localized around one representative path.
Effective metric and connection	(27.3), (27.4)	The projected quadratic form becomes curved kinematics with a gauge-like one-form.
Boundary equation of motion	(27.8)	The realized trajectory is geodesic-like with projected curvature and scalar forces.
Operational least-time form	(27.9)	The boundary selector is a generalized PLT.

27.1 The regime, and a worked parallel

The trajectory gauge holds where its slice is honest: the packet must stay sharply localized, which is the small-noise, high-visibility regime — diffusion weak against drift on the scale of the

path’s curvature. That is a statement, not an apology; it is the same regime condition under which Fermat’s rays emerge from wave optics and the classical path emerges from Feynman’s sum, now derived for an open system rather than assumed by analogy. Where the packet spreads, the trajectory language simply stops referring to anything, and selection is carried by the current gauge on the same functional. *What would refute the reading in its regime:* a selected path in a medium of independently measured mobility that departs systematically from the geodesic of the metric (27.3), beyond the small-noise corrections the functional itself predicts.

The oldest and the newest instances of the gauge are worth setting side by side, because they are the same calculation.

Table 19: The trajectory gauge run twice: Fermat’s ray of 1662 and a chemotactic run, column by column. Each row is one element of the derivation, and the two columns differ only in where the effective index comes from — prescribed by the glass in one case, computed from the projected mobility and drift in the other. The point of the table is that the right-hand column is not a metaphor: the bacterium’s path extremizes the same generalized time functional (27.9), in a medium whose geometry the projection derives rather than the modeler declares, and the regime condition in the last row is the honest price of the correspondence in both centuries.

Element	Light in graded glass (1662)	Chemotactic run (this Part)
Medium	Refractive index $n(z)$	Projected metric $g_{ab} = \frac{1}{2}D_{ab}^{-1}$ and drift one-form a_a
Where it comes from	Composition of the glass, measured	The projection: mobility and lawful drift of the packet, derived
Selected path	Snell’s law; ray bending toward high index	Geodesic of (27.8), bent by curvature and drift toward high effective index
Extremized quantity	Travel time $\int n \, ds$	Generalized optical time (27.9)
Regime	Wavelength small against the index gradient	Noise small against drift (sharp packet); visibility high
What breaks it	Diffraction	Packet spreading; selection passes to the current gauge

28 Current gauge: derivation of MEPP

The second gauge is the current gauge. Here the density ρ remains spatially extended and the fundamental question is no longer “which path does one packet follow?” but “which flux pattern is selected under the projected driving and constraints?” This is the natural language for transport networks, reaction networks, metabolic fluxes, ecological throughputs, and other non-equilibrium open systems.

Define the mobility or Onsager tensor

$$L^{ab}(z, t) := \rho(z, t) D^{ab}(z, t), \quad (28.1)$$

which is symmetric and positive semidefinite because D^{ab} is symmetric positive and $\rho \geq 0$. Next introduce the generalized force

$$X_a(z, t) := (D^{-1})_{ab} u^b(z, t) - \partial_a \varphi(z, t), \quad (28.2)$$

where φ is the scalar contribution extracted from the projected potential and any conservation constraints. The local entropy-production rate density is then written in the familiar bilinear form

$$\dot{\sigma}(z, t) := J^a X_a. \quad (28.3)$$

The total entropy-production rate is

$$\dot{\Sigma}[J] = \int_{\mathcal{X}} J^a X_a \, dz. \quad (28.4)$$

The same quadratic term that appeared in Eq. (26.6) may now be rewritten as an Onsager–Rayleighian,

$$\mathcal{R}_\rho[J; X] = \int_{\mathcal{X}} \left(J^a X_a - \frac{1}{2} J^a (L^{-1})_{ab} J^b \right) dz, \quad (28.5)$$

again subject to the continuity equation and the integrated Qi-QUA ledger constraints. The advantage of Eq. (28.5) is that its meaning is immediate: the first term rewards entropy-producing throughput under the imposed force field, whereas the second term penalizes dissipative expenditure measured by the inverse mobility tensor.

Varying Eq. (28.5) with respect to J gives

$$0 = \frac{\delta \mathcal{R}_\rho}{\delta J^a} = X_a - (L^{-1})_{ab} J^b, \quad (28.6)$$

hence

$$J^a = L^{ab} X_b. \quad (28.7)$$

Equation (28.7) is the force-flux law selected by the projected geometry. Substituting it back into Eqs. (28.4) and (28.5) yields

$$\dot{\Sigma}_\star = \int_{\mathcal{X}} X_a L^{ab} X_b \, dz \geq 0, \quad (28.8)$$

and

$$\mathcal{R}_{\rho, \star} = \frac{1}{2} \dot{\Sigma}_\star. \quad (28.9)$$

The second variation is

$$\delta^2 \mathcal{R}_\rho = - \int_{\mathcal{X}} \delta J^a (L^{-1})_{ab} \delta J^b \, dz \leq 0, \quad (28.10)$$

so the stationary point is a maximum of $\mathcal{R}_\rho$ on the admissible force-controlled slice. Because Eq. (28.9) identifies the optimum Rayleighian with one-half of the selected entropy-production rate, the same stationary point is also the maximum-entropy-production point on that slice.

This is the precise sense in which MEPP appears in Qi-EP. One must keep the control convention explicit. On the force-controlled slice, Eq. (28.5) is maximized. On the dual flux-controlled slice,

the Legendre transform

$$\mathcal{D}_\rho[J; X] = \frac{1}{2} \int_{\mathcal{X}} J^a (L^{-1})_{ab} J^b dz - \int_{\mathcal{X}} J^a X_a dz \quad (28.11)$$

is minimized. Thus the familiar “maximum entropy production” and “minimum dissipation” formulations are not contradictory in the present framework. They are dual descriptions of the same stationary point under dual control conventions (Dewar, 2003, Nicolis and Prigogine, 1977, Onsager, 1931, Prigogine, 1967, Seifert, 2012).

Theorem 28.1 (MEPP as the current gauge of Qi-EP(T-Qi4-10.4)). *On the force-controlled current slice $\mathcal{S}_{\text{flx}} \subset \Gamma_{\text{adm}}$, the projected action (26.6) is equivalent to the Onsager–Rayleighian (28.5). Its stationary point satisfies the flux-force law (28.7) and maximizes the entropy-production rate (28.4) on the admissible slice. The Legendre-dual flux-controlled description is Eq. (28.11), which is minimized at the same stationary point.*

Proof. Rewrite the quadratic term of Eq. (26.6) using $L^{ab} = \rho D^{ab}$. Up to terms independent of J , this gives Eq. (28.5). The first variation yields Eq. (28.7). The second variation (28.10) is negative semidefinite because L^{-1} is positive semidefinite on the support of ρ . Therefore the stationary point is a maximum of $\mathcal{R}_\rho$. At that point the identities (28.8) and (28.9) hold by substitution. The Legendre-dual representation follows by changing the control variables from X to J , yielding Eq. (28.11). ■

The duality just used deserves to be stated more sharply, because the second reading is not merely “minimum dissipation” in a loose sense: it is, exactly, the minimum-entropy-production theorem of Prigogine (1967). The two readings differ in what is held fixed, and each answers a different question.

On the *flux side* one clamps the forces and asks which constitutive relation nature adopts. Among fluxes compatible with power balance, the physical flux is the one of *greatest* entropy production, and it satisfies Ziegler’s orthogonality condition with a multiplier that the constraint fixes at $\lambda = \frac{1}{2}$ — not chosen, computed. This reading *selects the transport law*.

On the *force side* one fixes the constitutive relation, clamps some forces and frees the rest, and asks which state the system settles into. The stationary point of the force-side entropy production over the free forces is exactly the state whose free-force fluxes vanish — the steady state — and the restricted Hessian is positive definite, so it is a *minimum*. This reading *selects the state*, and it is Prigogine’s theorem.

Theorem 28.2 (MEPP and MinEP as conjugate readings of one stationary point (T-Qi4-10.7)). *Let $\dot{\Sigma}(J) = J^a (L^{-1})_{ab} J^b$ be the flux-side entropy production and $P(X) = X_a L^{ab} X_b$ the force-side one, with $L = L^\top \succ 0$ the mobility of Eq. (28.1). (i) Flux side, forces clamped. Among fluxes satisfying the power-balance constraint $\dot{\Sigma}(J) = J^a X_a$, the physical flux $J = LX$ of Eq. (28.7) maximizes $\dot{\Sigma}$, and the maximizer satisfies $X_a = \lambda \partial \dot{\Sigma} / \partial J^a$ with $\lambda = \frac{1}{2}$ fixed by the constraint. (ii) Force side, constitutive law fixed. Split the forces into clamped $X_1, \dots, X_m$ and free $X_{m+1}, \dots, X_n$. On the affine subspace of clamped forces, $\partial P / \partial X_j = 2L_{jk} X_k = 2J_j$ for free j , so the stationary point of P is the state with vanishing free-force fluxes, and the Hessian $2L_{jj'}$ on the free block is positive*

definite: the stationary point is a minimum. (iii) The two are Legendre-dual descriptions of the single stationary point of $\mathcal{R}_\rho$, and $\dot{\Sigma}(J_\star) = P(X_\star)$ there.

Proof. Given in full in Appendix F.6. ■

Remark 28.3 (What the duality settles). The two entropy-production principles have been presented as rivals for seventy years. Theorem 28.2 says they are not rival claims about nature at all: they are conjugate conventions about what is held fixed at one stationary point, one selecting the transport law and the other the state. The dissolution is not a diplomatic compromise between two schools; it is what the projected functional says when one asks it the two different questions, and it agrees with the resolution that the critical literature reached from below by direct circuit analysis (Bruers et al., 2007, Jaynes, 1980, Polettini, 2013). Two routes to the same reading, one from the projection and one from below, is a check rather than a coincidence.

28.1 The third reading, and Prigogine’s theorem recovered

Two identifications close the loop between this section and the history of Sections 22–23, and both are identities rather than analogies.

Onsager’s least dissipation is the Rayleighian itself. The functional (28.5) is, term for term, the 1931 functional $\sigma(J, X) - \Phi(J, J)$ of Section 22: the reward term is the entropy production and the penalty is the quadratic dissipation function with $R = L^{-1}$. Its maximization at fixed forces *is* Onsager’s principle of the least dissipation of energy, and its stationarity condition *is* the constitutive law (28.7). What the projection adds to 1931 is provenance: $L = \rho D$ is not a phenomenological matrix but the projected mobility, symmetric and positive because the bulk geometry is, so reciprocity arrives here as an inheritance rather than a separate microscopic argument. The three classical statements are thereby held in one hand: least dissipation encodes the *law*, the flux-side maximum selects the *transport* against power balance, and the force-side minimum selects the *state* — three questions, one stationary point.

Corollary 28.4 (Prigogine’s theorem, recovered with its hypotheses). *Restrict the force-side reading of Theorem 28.2(ii) to a packet whose projected mobility $L = \rho D$ is constant over the occupied region, whose boundary variables are even under time reversal, and whose ports clamp the forces $X_1, \dots, X_m$ while leaving the rest free. Then the stationary state minimizes the force-side entropy production $P(X)$ over the free forces, the minimum is the state with vanishing free-force fluxes, and P decreases monotonically along the relaxation. This is Theorem 22.1, hypothesis for hypothesis: (A1) linearity is the quadratic restriction of the projected action, (A2) constancy is the state-independence of the projected mobility, (A3) reciprocity is the symmetry of D , (A4) positive definiteness is its positivity on the packet’s support, and (A5) the clamped/free split is the ports declaration of SIP. Where any hypothesis fails, the corollary fails exactly as the 1945 theorem does — and Theorem 30.1 computes the failure instead of discovering it.*

The corollary is the checkpoint of Section 22 passed: the one proved theorem of the entropy-production tradition reappears, hypotheses and all, as the gentle-driving shadow of the projected functional. The seventy-year quarrel can now be closed in a table.

Table 20: The extremal statements of the entropy-production tradition, each placed at the one stationary point of the Rayleighian (28.5). Each row names the historical claim, what it varies, what it holds fixed, and what it selects; the last column locates the statement in the present derivation. The table is the settlement of the quarrel reviewed in Sections 22–23: the five claims were never rivals about nature, but five control conventions at one stationary point — except the last, which claimed what no convention delivers and failed as Section 23 records.

Claim	Varies	Holds fixed	Selects	Here
Onsager 1931, least dissipation	Fluxes	Forces	The transport law $J = LX$	$\mathcal{R}_\rho$ itself, Eq. (28.5)
Ziegler; Kohler–Ziman (MEPP, linear)	Fluxes on the power-balance surface	Forces	The transport law, as a maximum	Theorem 28.2(i), $\lambda = \frac{1}{2}$ computed
Prigogine 1945 (MinEP)	Free forces	Clamped forces; the law	The steady state	Theorem 28.2(ii); Corollary 28.4
Paltridge (climate MEPP)	Material transport	Radiative ports	The flux pattern	Flux-side reading with ports declared (§31)
Dewar 2003–2005 (general MEPP)	Path measure	Mean fluxes	A universal maximum — claimed	Not recovered; fails where Theorem 30.1 says it must

This current-gauge derivation is especially important for chemistry and life. A living system is not well described as a single sharply localized path. It is better described as a network of sustained throughputs, coupled channels, homeostatic gradients, and current constraints. In precisely that regime the most natural boundary language is MEPP or its dual minimum-dissipation form. The underlying selector, however, is still the same projected GG-6/GDOS variational principle.

Table 21: Step-by-step structure of the maximum-entropy-production derivation. Seven steps run down the first column, from the mobility tensor through the generalized force and the entropy production to the stationary flux law and the two readings of its dual. The layout deliberately matches Tables 18 and 23: the three chains start from the same projected functional and diverge only in the choice of primary variable, which is the content of the pillar’s common-origin theorem.

Step	Equation	Meaning
Mobility tensor	(28.1)	The projected metric becomes an Onsager transport tensor weighted by the density.
Generalized force	(28.2)	The projected drift and potential define the thermodynamic driving field.
Entropy production	(28.4)	The bilinear force-flux pairing is the boundary entropy-production rate.
Rayleighian	(28.5)	Entropy production competes with dissipative cost.
Stationary flux law	(28.7)	Extremization selects the physical current pattern.
Dual description	(28.11)	The same stationary point is Onsager’s least dissipation in the flux-controlled picture.
Force-side reading	Theorem 28.2(ii)	Holding the transport law fixed and freeing some forces, the same point is Prigogine’s minimum entropy production, selecting the state.

29 Inferential gauge: derivation of FEP

The third gauge is the inferential gauge. It becomes natural whenever the open subsystem does not merely transmit or dissipate, but also maintains internal states that summarize hidden causes and guide action. Cells, nervous systems, adaptive agents, and cognitive architectures all fall into this class.

Let the projected boundary variables be decomposed as

$$z = (\eta, s, a, \mu), \quad (29.1)$$

where η denotes hidden external causes, s sensory states, a active states, and μ internal states. The Markov-blanket interpretation is that the subsystem accesses the external world only through s , changes it through a , and maintains a compact internal model through μ (Friston, 2010). In the present framework the crucial addition is that the generative structure is not arbitrary: it is induced from the projected GG-6/GDOS geometry.

Let $p(\eta, s, a)$ denote the projected generative density inherited from the boundary stationary measure. Introduce a parametric recognition family $q_\mu(\eta)$ representing the subsystem’s internal approximation to the posterior over hidden causes. The variational free energy is then

$$\mathcal{F}(s, a, \mu) = \int q_\mu(\eta) \log \frac{q_\mu(\eta)}{p(\eta, s, a)} d\eta. \quad (29.2)$$

As usual, Eq. (29.2) may be rewritten as

$$\mathcal{F}(s, a, \mu) = D_{\text{KL}}(q_\mu(\eta) \parallel p(\eta|s, a)) - \log p(s, a), \quad (29.3)$$

so $\mathcal{F}$ upper-bounds surprisal and is minimized when the recognition family matches the posterior.

To recover FEP from the projected action, restrict the density ρ to the statistical manifold

$$\mathcal{S}_{\text{inf}} = \{\rho_{\mu, s, a}(\eta) = q_\mu(\eta) \delta(s - s_t) \delta(a - a_t)\}. \quad (29.4)$$

When the quadratic term of Eq. (26.6) is pulled back to this manifold, it induces the Fisher metric

$$g_{ij}^{\text{F}}(\mu) = \int q_\mu(\eta) \partial_i \log q_\mu(\eta) \partial_j \log q_\mu(\eta) d\eta, \quad (29.5)$$

which is the natural information-geometric metric on the parameter manifold (Amari, 1998, Amari and Nagaoka, 2000, Jaynes, 1957). The restricted action becomes

$$\mathcal{A}_{\text{inf}}[\mu] = \int_0^T dt \left[\frac{1}{2} g_{ij}^{\text{F}}(\mu) \dot{\mu}^i \dot{\mu}^j + \mathcal{F}(s(t), a(t), \mu(t)) \right]. \quad (29.6)$$

The Euler–Lagrange equations are

$$g_{ij}^{\text{F}} \ddot{\mu}^j + \Gamma_{ijk}^{\text{F}} \dot{\mu}^j \dot{\mu}^k + \partial_i \mathcal{F} = 0, \quad (29.7)$$

where Γ_{ijk}^{F} is the Levi–Civita connection of the Fisher metric. In the overdamped or adiabatically relaxed limit — the relevant limit for slow inferential updates compared with fast hidden fluctuations — the inertial term may be neglected, leaving

$$\dot{\mu}^i = -\kappa_\mu (g^{\text{F}})^{ij} \partial_j \mathcal{F}. \quad (29.8)$$

Equation (29.8) is the natural-gradient version of the FEP: internal states descend variational free energy in the Fisher geometry. Active states obey the analogous projected descent

$$\dot{a}^\alpha = -\kappa_a \frac{\partial \mathcal{F}}{\partial a^\alpha}, \quad (29.9)$$

which is the action side of active inference: the system acts on the environment to reduce sensory prediction error.

Theorem 29.1 (FEP as the inferential gauge of Qi-EP(T-Qi4-10.5)). *On the statistical slice $\mathcal{S}_{\text{inf}} \subset \Gamma_{\text{adm}}$, the projected action (26.6) pulls back to the Fisher-metric action (29.6). In the overdamped limit, its stationary dynamics reduce to the natural-gradient FEP equations (29.8) and (29.9). Hence FEP is the inferential gauge of Qi-EP.*

Proof. The pullback of the density family $\rho_{\mu, s, a}$ onto the parameter manifold μ converts the quadratic kinetic term of Eq. (26.6) into the Fisher quadratic form (29.5); this is derived explicitly in Appendix F.3. The projected potential becomes the variational free energy $\mathcal{F}$, because the hidden-sector elimination identifies the admissible statistical mismatch cost with the KL divergence

between q_μ and the projected generative density. The Euler–Lagrange variation of Eq. (29.6) yields Eq. (29.7). The overdamped limit removes the second-order term and leaves the natural-gradient descent (29.8). The active-state update (29.9) follows by the same reduction applied to the action variables. ■

This derivation is the formal bridge from the present physics paper to later life, brain, and mind applications. It says that a persistent adaptive subsystem can be described as an open-system packet living on a statistical manifold whose geometry is inherited from the same projected bulk action. The resulting inferential descent is not a separate law introduced only for neuroscience. It is the inferential gauge of the same Qi-EP selector that also appears as PLT and MEPP.

29.1 What the inferential gauge does not assume

The inference reading is the one most often overclaimed, so it is worth stating precisely which of its usual supports the present derivation does *not* lean on. Four points, each of which the framework replaces with something checkable.

The blanket is replaced by declared ports. Nothing above assumes that a statistical boundary can be found inside the stationary density of a generic system. The decomposition (29.1) is read here as *constitutive data*: the ports of the system packet SIP, declared when the system is written down and carrying the criteria under which the description is judged wrong (§4.3). This is the same move by which thermodynamics escapes the ambiguity of “the system” — by declaring its walls rather than hoping to discover them (Aguilera et al., 2022).

The synchronization map is replaced by a rank condition. That the internal states parameterize a posterior over external ones is not assumed here. The internality construction of Part IX makes the internal model’s existence conditional and checkable: the external perturbation is recoverable from the sensory history if and only if an explicit action-conditioned observability matrix has full column rank, and where the rank fails, no unique internal model exists. Systems that do not satisfy the inference reading are thereby classified rather than assumed away.

The obstruction is the one already located. The couplings that must be excluded for the inference reading to hold are, in boundary coordinates, curl-sector couplings — the same $\mathcal{C}_{\text{curl}}$ of Theorem 30.1. The inferential gauge holds exactly where the current gauge’s extremum readings hold, and fails with them, for the same reason. One obstruction, three literatures.

What is not derived is not claimed. The expected-free-energy functional of prospective policy selection does not follow from Eq. (29.6). In the present framework prospective selection is carried by a different mechanism — the internalized regime switch of the runtime — and no claim is made that Eq. (29.8) contains it.

The four replacements are worth fixing in a table, because they are the exact points at which the critical literature of Section 24 located the free energy principle’s troubles — and the derivation must be seen to answer each one rather than to relabel it.

Table 22: The inferential gauge against the free energy principle it recovers, assumption by assumption. Each row is one structure the FEP literature supplies by hand, with the critique that exposed it (Section 24) and the derived object that replaces it here. The last row is the one no formulation of the principle has ever supplied for itself: a stated regime, with the failure outside it carried by the same theorem that bounds the other two gauges. The table is the section’s claim in compact form: the working mathematics of the tradition transfers unchanged, while every contested structure arrives derived, conditional, or explicitly not claimed.

Structure	In the FEP literature	In the inferential gauge
Boundary	Markov blanket, sought in the stationary density; definitions inequivalent (Biehl et al., 2021)	Ports of the packet, declared as constitutive data (SIP)
Metric on beliefs	Chosen with the model class	Fisher metric, pulled back from the projected quadratic form (29.5)
Internal model exists	Synchronization map, assumed; non-generic (Aguilera et al., 2022)	Rank condition of the internality construction; checkable, and classifying where it fails
Solenoidal coupling	Excluded silently (Friston et al., 2021)	The curl sector of Theorem 30.1: the exclusion is the stated regime
Policy selection	Expected free energy, added	Not claimed; carried by the runtime’s regime switch
Regime of validity	Not stated; contested	Stated: the lawful sector, failing with the other gauges by Theorem 30.1

Table 23: Step-by-step structure of the free-energy-principle derivation. Six steps run down the first column, from the blanket decomposition through the projected generative model and the Fisher metric to the inferential action and the natural-gradient descent it induces. This is the chain that matters most outside physics: it obtains the free-energy principle used across theoretical neuroscience as a *gauge* of the same functional that gives least time and maximum entropy production, rather than as a separate principle introduced for living systems.

Step	Equation	Meaning
Blanket decomposition	(29.1)	External hidden causes, sensory states, active states, and internal states are separated.
Projected generative model	(29.2)	The subsystem compares its recognition density with the projected world-model.
Statistical slice	(29.4)	The current-density action is restricted to a parametric statistical manifold.
Fisher metric	(29.5)	The projected kinetic term becomes information geometry.
Inferential action	(29.6)	The same reduced action becomes a Fisher-kinetic plus free-energy functional.
Natural-gradient descent	(29.8)	In the overdamped limit the selector becomes FEP.

30 Where the three gauges must fail

A derivation that produced only successes would be worth less than one that produces its own limits. This subsection computes exactly where the three gauges stop holding, and the answer is

already written in the runtime’s own bookkeeping: in the two sectors of the boundary current that neither descend the ledger nor diffuse.

Recall the four-sector decomposition (4.4) — gradient, curl, noise, and gauge — and the balance identity (8.3) it induces on the ledger, $d\mathcal{F}_{\text{GFF}}/dt = -\mathcal{D}_{\text{GFF}} + \mathcal{P}_{\text{port}} + \mathcal{C}_{\text{curl}} + \mathcal{G}_{\text{gauge}}$. The reduction of the projected action (26.6) to the Onsager–Rayleighian (28.5) used two properties of the lawful drift, both of them exactly the conditions under which the last two terms vanish: the curl current is tangent to the ledger’s level sets, so circulation does no work against the ledger, and the ledger is gauge-invariant, so the gauge sector does none either. Where either condition fails, the reduction loses a term — and with it, the whole entropy-production family loses its stationarity.

Theorem 30.1 (Where the extremum readings must fail (T-Qi4-10.8)). *Let the lawful drift be split as $\rho u = \rho u_{\text{d}} + \rho u_{\text{c}} + \rho u_{\text{g}}$, with u_{d} collecting the gradient and noise sectors of (4.4). Then the physical stationary current of the full projected action is*

$$J_{\star}^a = L^{ab} X_b + \rho u_{\text{c}}^a + \rho u_{\text{g}}^a, \quad (30.1)$$

and the first variation of $\dot{\Sigma}$ at $J_{\star}$ retains a term linear in the variation whose integral against the potential-generated family of admissible variations is $-(\mathcal{C}_{\text{curl}} + \mathcal{G}_{\text{gauge}})/\Theta_{\text{eff}}$. Consequently, wherever $\mathcal{C}_{\text{curl}} \neq 0$ or $\mathcal{G}_{\text{gauge}} \neq 0$, no functional of the entropy-production family is stationary at the physical state: the flux-side and force-side readings of Theorem 28.2 fail together, and the inferential gauge fails with them.

Proof. Given in full in Appendix F.7. ■

The theorem does not report that the extremum readings sometimes fail. It computes the exact tax that makes them fail. A term linear in the variation survives at the physical state, and a point carrying a surviving linear term sits on a slope, not on a summit. That is the whole mechanism, and the classical counterexamples are three names for that one slope. Figure 9 draws both cases on one baseline.

Remark 30.2 (The counterexamples, located). The known failures of the entropy-production principles land in the two obstructing sectors and nowhere else. A driven circuit with an inductive branch fails the force-side reading because an inductive current is odd under time reversal: in boundary coordinates it is curl-sector circulation doing work, $\mathcal{C}_{\text{curl}} \neq 0$. A network held at two temperatures fails because a non-uniform temperature field makes the ledger’s level sets and the circulation non-orthogonal, so the work terms of (8.3) survive (Jaynes, 1980). And the modern statement that selection beyond linear order depends on time-symmetric activity that no entropy-production functional contains is the observation that $\mathcal{C}_{\text{curl}}$ and $\mathcal{G}_{\text{gauge}}$ are time-symmetric sector couplings invisible to $\dot{\Sigma}$ (Bruers et al., 2007, Poletti, 2013). The framework does not rescue the extremum readings in these cases; it derives that they must fail there, which is what a correct derivation owes.

Remark 30.3 (Beyond quadratic order, and why the confinement is a theorem). The Onsager–Rayleighian (28.5) is the quadratic restriction of the projected action (26.6). The full functional is not quadratic, and its stationarity conditions beyond quadratic order are not expressible as an extremum of any entropy-production functional — in agreement with, and for the same structural

reason as, the macroscopic-fluctuation-theory result that what survives far from equilibrium is a corrected minimum-cost principle for the full functional rather than an entropy extremum (Bertini et al., 2015). The confinement of both entropy-production readings to the linear regime is therefore a theorem about the projection, not an empirical accident of the systems in which they were first tried.

The three results of this subsection — the duality, the failure location, and the order confinement — are what convert Qi-EP from a unification into a falsifiable claim. A framework that derived the three principles and nothing else would be a restatement; one that also derives, from the same object and with no new input, the exact conditions under which each must break is making a claim that measurement can refute.




Figure 9: The mechanism of Theorem 30.1, drawn. Both panels plot the entropy production $\dot{\Sigma}$ against an admissible variation δ about the physical stationary current J_* of Eq. (30.1); the curvature is the same in both, because it comes from the same positive-definite inverse mobility. On the left the curl and gauge sectors do no work against the ledger, the cross term of the reduction vanishes, and J_* sits at the summit — which is what makes the flux-side and force-side readings of Theorem 28.2 extremal characterizations at all. On the right one of those sectors does work, a term linear in δ survives at J_* , and the same curve is tilted: the dashed tangent no longer lies flat, and a point with a surviving linear term sits on a slope rather than on a summit. The tilt is not qualitative — it is $-(C_{\text{curl}} + G_{\text{gauge}})/\Theta_{\text{eff}}$, so the work the two sectors do against the ledger *is* the slope, and no functional of the entropy-production family is stationary where that work is nonzero. Read the two panels as one statement: the extremum readings are not approximately true outside the lawful sector, they are structurally absent there.

31 The three gauges in the laboratory

The derivation would be idle if the gauges did not read real systems, and unfalsifiable if each reading did not name the measurement that would break it. This is where the argument meets things one can point at: a thermodiffusion cell, a resistor network, a pan of convecting fluid, three planetary atmospheres, the enzyme that makes ATP, a swimming bacterium, and a predicting brain. The format is the same for every case and deliberately blunt — what is observed, how the old principle read it, how the derived gauge reads it, and the measurement that would prove the

new reading wrong. None of the eight is offered as a confirmation; they are placements, and the discipline of the register is to say which rows could bite.

31.1 Path selection in a structured medium

Observed. A particle or excitation crossing a medium of varying mobility selects one path among many — the situation the least-time law was historically built on.

The old reading. Fermat’s principle, exact for rays, silent on noise.

The gauge reading. The trajectory gauge in its home regime: a least-time law with the effective metric fixed by the projected mobility rather than declared, per Table 19.

What would refute it. A selected path in a medium of independently measured mobility that departs systematically from the geodesic of that metric, beyond the small-noise corrections the functional itself predicts.

31.2 Thermodiffusion: the textbook force-side state

Observed. A fluid mixture held between plates at different temperatures develops a stationary concentration gradient — the Soret effect — carrying heat but no net matter flow.

The old reading. Prigogine’s theorem in its home regime: with the thermal force clamped and the diffusive force free, the stationary state minimizes entropy production, realized by the vanishing of the matter flux.

The gauge reading. The force-side reading of Theorem 28.2(ii), exactly as in the classical proof; the derivation adds only the provenance of the quadratic form — the ledger’s Hessian — and the statement of scope: constant coefficients across the cell, all variables even under time reversal.

What would refute it. A measured stationary Soret state carrying a persistent matter flux at small driving, or a minimum displaced from the zero-flux state beyond the quadratic corrections.

31.3 Planetary heat transport: the flux side at scale

Observed. With no fluid dynamics — no rotation rate, no tuned diffusivity — an energy balance closed by maximizing the material entropy production of the meridional transport reproduces the equator-to-pole contrast on Earth, Mars, and Titan (Section 23).

The old reading. The flagship success of MEPP, with its standing embarrassments: nothing in the principle says which entropy production to maximize, and the choice that works ignores the radiatively dominant part of the budget.

The gauge reading. The flux-side reading with the radiative channels entering as clamped ports and the material transport as varied current — which is what makes an energy balance fed only by sunlight reproduce the contrast, and which converts Paltridge’s empirical constraint choice into the declared ports of the packet. The ledger, fixed upstream, identifies the material entropy production as the extremized quantity; what was arbitrary becomes stated, and therefore falsifiable.

What would refute it. An atmosphere, observed or adequately resolved in simulation, whose material transport sits systematically off the flux-side stationary point with radiative ports genuinely clamped and scale separation genuinely present — or a dependence of the selected state on resolution

that survives the stated scale-separation condition. Both concerns are live in the climate literature, and the reading stands exposed to them.

31.4 Jaynes’s two-temperature network: the counterexample as a theorem

Observed. A resistor network with elements held at different temperatures distributes its currents by Kirchhoff’s laws — minimizing total heat production, not total entropy production; the two differ whenever the temperatures do (Section 22).

The old reading. A refutation of MinEP’s generality — and it is.

The gauge reading. Theorem 30.1 with the inhomogeneity mechanism: a temperature field varying across the system makes the ledger’s level sets and the circulation non-orthogonal, the work terms survive, and the extremum principle is *derived* to fail here. The physical selection remains the stationary point of the full projected cost, which does not degenerate to an entropy functional in this case.

What would refute it. A driven network whose physical current distribution deviated at leading order from the stationary point of the full quadratic cost — not merely from the entropy extremum.

31.5 Bénard convection: the boundary of both readings

Observed. A fluid layer heated from below conducts at small driving and, above a critical threshold, organizes into convection rolls whose transport exceeds conduction.

The old readings. For MinEP, the historical falsifier: the convective branch produces more entropy and it is the convective branch that is stable. For the maximum tradition, the historical motivation.

The gauge reading. A regime boundary, not a paradox: below threshold the curl sector does no ledger work and the force-side minimum governs; the instability is the point at which a circulating sector begins to do work, the quadratic reduction fails per Theorem 30.1, and selection passes to the full functional. The framework does not predict the convective pattern from an entropy extremum — and says why no such prediction can be exact.

What would refute it. A driven steady state beyond threshold demonstrably selected by an entropy-production extremum alone, robustly across geometries — which the failure theorem forbids.

31.6 Chemotaxis: the inferential gauge with a checkable rank condition

Observed. A swimming bacterium climbs a chemical gradient it cannot spatially resolve, comparing present against remembered concentration and modulating its tumbling rate.

The old reading. The free energy literature’s minimal agent: an internal state inferring an external one across a blanket.

The gauge reading. The natural-gradient law (29.8) with a one-dimensional belief coordinate — the methylation state as the estimate of gradient sign — and the rank condition doing real work: the internal model exists *because* self-motion raises the observability rank, and a tethered cell, sampling but not moving, falls below the condition and can hold no unique gradient estimate. Receptors and

motor are the declared ports; the existence of the inference reading is a checkable condition on the sensorimotor loop rather than an assumption about the cell.

What would refute it. Gradient-climbing performance surviving closure of the action channel — immotile cells inferring gradients they cannot move through — or an identified chemotactic architecture whose update law departs from natural-gradient structure where the metric is measurable.

31.7 Neural prediction: the gauge the principle was built for

Observed. Cortical responses attenuate to predicted input and amplify prediction error; hierarchical processing carries top-down predictions against bottom-up residuals; behavioral and neural data are fit, quantitatively, by precision-weighted prediction-error dynamics (Section 24).

The old reading. The free energy principle’s home territory — where the working mathematics was always solid regardless of the foundational claims.

The gauge reading. The inferential gauge in its hierarchical Gaussian restriction, where (29.8) reduces to exactly the predictive-coding dynamics of the process theories — so everything that literature has established empirically transfers unchanged. What the derivation adds is the boundary of the reading (strongly circulating dynamics across the declared ports leave the gradient-flow regime) and a bridge that is checkable in principle: the metric is fixed by the ledger, so precision weighting becomes a measurable geometry rather than a free parameter family.

What would refute it. Identified neural inference dynamics whose metric structure, where measurable, is not Fisher — or inference maintained where the circulating sector demonstrably does ledger work, without the degradation the reduction predicts.

31.8 ATP synthesis: all three gauges in one molecule

Observed. The rotary enzyme that synthesizes ATP pays for each bond with an integer number of proton crossings set by its rotor size, while the bond’s free energy is universal — the measured standard hydrolysis free energy against the derived $2Q_B$ of Qi-QUA, at the percent level (§17).

The old readings. Bioenergetics treats the number as measured; the free energy literature treats the cell as an inference engine; the thermodynamic literature treats the chemiosmotic machine as a near-equilibrium transducer.

The gauge reading. The one system in this register that exercises all three gauges at once. The proton path through the rotor is trajectory-gauge — a least-time path through a curved potential. The steady turnover is current-gauge — flux–force operation of a linear transducer, with the duality of Theorem 28.2 fixing the operating point. The regulatory network around it is inferential-gauge — the cell’s estimate of its energy state updated by natural-gradient descent. The three descriptions do not conflict because they are three restrictions of the one projected functional (26.6) — which is the concrete answer to the question this Part opened with, of why a cell obeying three principles is not pulled in three directions.

What would refute it. An elementary catalytic event whose four-way bookkeeping violates the single conversion constant outside the stated tolerance — the falsifier of Qi-QUA, inherited here.

31.9 What the eight cases share

In every case the derived reading does three things the original principle could not: it says *which* functional, fixed upstream and never chosen per application; it says *where the reading stops*, as a stated sector or rank condition rather than a surprise; and it names a measurement that could kill it. And in every case the failure literature of Sections 22–24 reappears not as embarrassment but as confirmation of a boundary the derivation drew in advance. That inversion — counterexamples as theorems — is the practical content of deriving a principle instead of adopting it.

Table 24: The pillar’s observational register: which gauge each case exercises, and the measurement that would refute the reading. The right-hand column is the one that matters — a derivation that unified three principles without naming what would break each reading would have bought coherence at the price of content. Two of the eight cases sit past the failure boundary of Theorem 30.1 and are listed as such: there the framework predicts that no entropy-production extremum can select the state, and a demonstration to the contrary refutes it.

Observation	Gauge	What would refute the reading
Path selection in a structured medium	Trajectory	A selected path departing systematically from the geodesic of the independently measured mobility metric
Thermodiffusion (Soret cell)	Current, force side	A stationary state carrying persistent matter flux at small driving
Planetary heat transport	Current, flux side	Material transport systematically off the flux-side stationary point with ports clamped and scales separated
Two-temperature driven network	Past the failure boundary	Selection by an entropy-production extremum where circulation does ledger work
Convection beyond threshold	Past the failure boundary	A driven steady state beyond threshold selected by an entropy extremum alone, robustly across geometries
Chemotaxis	Inferential	Gradient climbing surviving closure of the action channel; or non-natural-gradient updates where the metric is measurable
Neural prediction	Inferential	Measured inference metric that is not Fisher; or inference without degradation where the circulating sector does ledger work
ATP synthesis	All three	Four-way event bookkeeping violating the single conversion constant beyond tolerance

The register is deliberately short and deliberately answerable. Five of the eight rows would refute a gauge reading, two would refute the failure theorem that bounds the gauges, and one — the last — would refute the pillar upstream of all of them. That distribution is the honest shape of a claim that unifies: most of its exposure is inherited, and the part that is its own is the part that says where it stops.

32 The unified Qi-EP theorem: one projected functional, three gauges

The three previous subsections are the heart of the section. They show that PLT, MEPP, and FEP are not independent principles living in separate disciplines. They are three reductions of one master current-density action.

It is useful to package this result formally. Define the quadratic form

$$\mathbb{K}_\rho(\delta J, \delta J) := \int_{\mathcal{X}} \delta J^a (\rho D)_{ab}^{-1} \delta J^b dz, \quad (32.1)$$

and write the projected action compactly as

$$\Omega_{\text{Qi-EP}}[\rho, J] = \int_0^T dt \left[\frac{1}{4} \mathbb{K}_\rho(J - \rho u, J - \rho u) + \mathcal{V}[\rho] \right] \quad \text{subject to} \quad \partial_t \rho + \nabla \cdot J = 0. \quad (32.2)$$

Then the three gauges may be summarized as follows:

$$\text{trajectory gauge: } (\rho, J) \mapsto (\delta_{z(t)}, \delta_{z(t)} \dot{z}(t)) \implies \Omega_{\text{Qi-EP}} \rightarrow \mathcal{T}[z], \quad (32.3)$$

$$\text{current gauge: } (\rho, J) \mapsto (\rho, J) \implies \Omega_{\text{Qi-EP}} \rightarrow \mathcal{R}_\rho[J; X], \quad (32.4)$$

$$\text{inferential gauge: } (\rho, J) \mapsto \rho_{\mu, s, a} \implies \Omega_{\text{Qi-EP}} \rightarrow \mathcal{A}_{\text{inf}}[\mu]. \quad (32.5)$$

Theorem 32.1 (Common-origin theorem for Qi-EP: three gauges of one projected functional (T-Qi4-10.6)). *Within the GG-6/GDOS framework, the projected functional (32.2) generates the three boundary principles*

$$\text{PLT,} \quad \text{MEPP,} \quad \text{FEP}$$

as its trajectory, current, and inferential gauges respectively. Therefore the three principles are gauge-equivalent boundary faces of one projected 6D variational law.

Proof. The proof is the composition of Theorems 27.1, 28.1, and 29.1. Each theorem establishes that a distinct admissible slice of Γ_{adm} together with a distinct coarse-grained coordinate choice reduces the same master functional (32.2) to one of the three familiar boundary forms. Since the starting functional is the same in all three cases, the resulting principles are not independent axioms but gauge-equivalent reductions. ■

The most important conceptual consequence is that Qi-EP restores unity across scientific domains. A transport problem, a metabolic steady state, and an adaptive inferential system no longer require three unrelated foundational principles. They require one projected selector viewed in three different gauges. This is exactly the kind of cross-domain unification that Paper IV-B is meant to achieve.

Table 25: The three gauges of Qi-EP. Each row is one gauge — trajectory, current and inferential — with its primary variable, the reduced functional it extremizes, and its operational interpretation. The table is the pillar’s claim in one page: least time, maximum entropy production and the free-energy principle are not three principles but one, read with three different choices of what is held primary, and the second column is where the choice is made.

Gauge	Primary variable	Reduced functional	Operational interpretation
Trajectory gauge	$z(t)$	$\mathcal{T}[z]$	Path or branch selection under an effective metric and connection (PLT).
Current gauge	(ρ, J)	$\mathcal{R}_\rho[J; X]$	Force-controlled flux selection and entropy-production extremization (MEPP).
Inferential gauge	$\mu(t)$	$\mathcal{A}_{\text{inf}}[\mu]$	Natural-gradient descent of variational free energy (FEP).

Table 26: The grade of every step of the Qi-EP chain, in the program’s standard grading. “Derived” means proved here or in the cited companion from the stated imports; “conditional” names the condition and states it; “argued” means supported but not proved; “not claimed” marks the adjacent results this pillar deliberately does not assert. The right-hand column is the load-bearing one: a pillar whose subject is three principles that were historically overclaimed must not add a fourth example, and this table is where the boundary of the present claim is drawn.

Statement	Grade	Condition / authority
Projected functional (26.6) from the 6D action	Derived	§26, via OGGEF
Trajectory gauge $\Rightarrow$ PLT	Derived	Theorem 27.1; small-noise limit
Current gauge $\Rightarrow$ flux-force law	Derived	Theorem 28.1
MEPP/MinEP duality at one stationary point	Derived	Theorem 28.2; Appendix F.6
Failure location of both readings	Derived	Theorem 30.1; Appendix F.7
Inferential gauge $\Rightarrow$ natural-gradient FEP; Fisher metric from CEP	Derived	Theorem 29.1; Appendix F.3
Internal-model existence	Conditional	rank condition of the internality construction (Part IX)
Statements inheriting the imported cascade closure	Conditional	GG-A6-GRIP’s conditioning, imported at its own grade (§3.5)
Exhaustiveness of the three gauges	Argued	from the packet’s structure (state, flow, registration); not proved
Expected free energy / policy selection from (29.6)	Not claimed	does not follow; carried by a different mechanism
Universal validity of MEPP	Not claimed	false outside its regime, and derived to be false (Theorem 30.1)

32.1 The contract, discharged

Section 25.1 fixed five criteria before the derivation began; the Part should not end without checking them off against what was actually proved. (i) *One derived object*: the three principles

were obtained from the single projected functional (32.2), itself the OGGEp projection of the bulk action rather than a fresh postulate (Proposition 26.1). (ii) *The restrictions as theorems*: the small-noise confinement of the trajectory reading (Theorem 27.1), the linear-regime confinement of both entropy-production readings (Remark 30.3), and the sector and rank conditions of the inferential reading all arrive derived, with the failure locations computed in Theorem 30.1 rather than discovered by counterexample. (iii) *The quarrel dissolved*: Theorem 28.2 and Table 20, in agreement with the resolution the critical literature reached from below. (iv) *Modeler-chosen structure replaced*: the ledger fixes which entropy production, the declared ports fix the boundary, and the rank condition replaces the synchronization map (Table 22). (v) *The conditional and open elements stated*: Table 26, where the grades are printed and nowhere waived. What a reader should retain is not that the three principles are true — each is true exactly where its regime theorem says — but that the three regimes, the three laws, and the three failures now follow from one object.

33 DGI, GSSB, and GA inside the Qi-EP landscape

Qi-EP is a projected selector, and a selector is not yet a dynamics: it says which configurations are stationary, not how a system moves among them, splits between them, or stays in one. Those three questions belong to the runtime engine of Theorem 7.3, and its three moves — DGI, GSSB, and GA — act on the Qi-EP landscape in three specific, checkable ways. This section records the action of each, because the three moves are also where the three gauges’ most familiar anomalies — caustics, instabilities, bistable percepts — turn out to be one phenomenon wearing three costumes.

DGI acts as the gauge orbits of the selector. Admissible informational-gauge transformations — relabelings of the boundary description that carry no ledger cost — change $\mathfrak{Q}_{\text{Qi-EP}}$ by a boundary term only, so stationary points come in gauge-equivalence classes, and the physical content of the selector lives on the orbit space. Each gauge shows the orbit in its own clothes. In the current gauge it is the gauge sector of the four-current (4.4): flow that relabels without transporting. In the inferential gauge it is reparameterization of the belief coordinates: the natural-gradient law (29.8) is form-invariant under $\mu \rightarrow \mu'(\mu)$ precisely because the Fisher metric transforms as a tensor — which is why the derivation produced the natural gradient rather than the naive one, and why “which coordinates do the beliefs use” is a question with no physical content.

GSSB acts as branch creation. Where the Hessian of the selector degenerates, $\det \partial^2 \mathfrak{Q}_{\text{Qi-EP}} = 0$, a stationary point splits, and the runtime commits the system to one branch of the split. The three gauges meet this one condition as three famous phenomena. In the trajectory gauge it is the kinetic focus: the point beyond which the extremal path becomes a saddle (Section 21) and neighboring extremals cross — the caustic of optics. In the current gauge it is the instability threshold: Bénard onset is the degeneration of the quadratic form as a circulating sector begins to do work, which is why the entropy readings lapse exactly there (§31). In the inferential gauge it is bistable perception: a posterior landscape whose single mode splits, defeating the single-mode assumption cataloged in Section 24. One degeneracy condition, three literatures’ edge cases — the same economy the failure theorem showed for the counterexamples.

GA acts as the persistence criterion. A stationary point that the runtime selects must also be one it can keep: in the gradient sectors the ledger descends monotonically ($\dot{\mathcal{F}}_{\text{GFF}} \leq 0$ along the lawful flow), so the selected extremum is an attractor rather than a transient, and the depth of its basin is the resource that Part VII will convert into a persistence class. This is the clause that connects the selector to everything downstream: the SeOS membership conditions of Part VII, and, in the companion capstone, the Ground Geometric Attractor whose basin-by-basin succession carries the eighteen-layer cascade (Arisaka, 2026, A11-Origin).




Figure 10: The three moves of the runtime engine acting on the Qi-EP landscape, drawn one per panel. Left, DGI: gauge transformations carry no ledger cost, so a stationary point is a ring of equivalent minima and the physical content lives on the orbit space — in the inferential gauge this is exactly why the natural gradient, and not the naive one, appears in Eq. (29.8). Center, GSSB: where the selector’s Hessian degenerates a valley splits, and the three gauges meet this single condition as the caustic of optics, the convective instability of the current reading, and the bistable percept of the inferential one — three literatures’ edge cases as one degeneracy. Right, GA: in the gradient sectors the ledger descends monotonically, so what the selector picks the runtime keeps; the basin’s depth is the quantity Part VII converts into a persistence class. The figure is a schematic of mechanisms, not a computed surface, and each panel’s claim is carried by the theorem it cites in the text.

34 Why Qi-EP applies to all open systems

At this stage it is worth stepping back and asking why the section matters beyond formal derivation. The answer is that every natural open system presents some combination of the three selector questions:

- (i) Which path or branch is followed?
- (ii) Which flux pattern or throughput is sustained?
- (iii) Which internal model or policy is stabilized?

A purely physical transport problem may use almost exclusively the first question. A reaction network, molecular motor, or ecosystem may use primarily the second. A cell, brain, or adaptive agent inevitably uses the third. But no natural system is forbidden from moving between gauges. A cell chemotaxes along a path (PLT language), maintains metabolic fluxes (MEPP language), and

updates internal states in response to sensory mismatch (FEP language). The same object therefore occupies all three gauges at once, depending on which variables have been retained by the observer.

This is why the present derivation cannot be restricted to one discipline. If Qi-EP were only a neuroscience principle, it could not serve as Pillar II of a grand open-system framework. If it were only a non-equilibrium thermodynamic principle, it could not later support Qi-HAL and the inner-hologram architecture. If it were only a geodesic-trajectory principle, it could not explain living systems. The point of the present derivation is exactly that the same projected 6D variational law has enough flexibility to cover all three uses without multiplying foundational assumptions.

The universality claim can be summarized in a domain table.

Table 27: Illustrative domain-wise readings of Qi-EP. Six domains run down the first column, from particle and wave transport to societies and collective systems. The second column names the gauge that is most natural there, the third the object it selects, and the fourth an illustrative reading. The word *illustrative* in the heading is load-bearing: these are organizational placements meant to orient a reader in an unfamiliar domain, not predictions, and none of them is offered as evidence.

Domain	Most natural gauge	Selected object	Illustrative reading
Particle and wave transport	PLT	Path or branch	Effective geodesic, tunneling channel, or propagation branch under the projected boundary geometry.
Chemistry and reaction networks	MEPP	Flux pattern	Stationary throughput through coupled channels under force-controlled or resource-controlled driving.
Molecular machines	MEPP + PLT	Cycle and substep path	Driven cycles and substep selections constrained by the Qi-QUA budget.
Cells and homeostatic systems	MEPP + FEP	Metabolic plus inferential state	Throughput must be maintained while internal states track the hidden environment.
Brains and cognition	FEP + PLT	Inference trajectory and policy	Prediction, sensing, and action unfold on an information-geometric manifold.
Societies and collective systems	All three	Path, flow, and policy	Collective trajectories, throughputs, and internal models are all needed simultaneously.

34.1 Nothing biological was used

The reach of the theorems is set by what their proofs consumed, and it is worth listing how short that list is. The derivation used a packet with declared ports, a ledger, and a projection — SIP, GFF, and OGGEF/GDOS — and at no step a membrane, a genome, a neuron, or any property that distinguishes living matter from driven matter. The scope of the three gauge theorems is therefore the class of systems carrying that structure: every open system admitting a packet description, and in particular the entire SeOS class of Part VII, whose membership conditions are stated there as

five checkable clauses and whose persistence theorem is what converts “selected” into “still there to be measured.” A living organism and a brain are members of that class by those theorems — not by analogy, not by aspiration, but by satisfying the clauses — and so the three principles reach organisms and minds with exactly the regimes and exceptions derived above, and with nothing further granted. This is the sense, and the only sense, in which the Part’s claim extends to life: the theorems bind the class, and life is in the class.

34.2 The cascade demonstration

The claim of the previous subsection is structural; the companion capstone makes it a demonstration (Arisaka, 2026, A11-Origin). There the GRIP-In/GRIP-Out engine, running under the Qi-4 admissibility conditions of the present paper, drives the eighteen-layer cascade of Ground Geometric Attractors from the gauge group and the electroweak vacuum through the periodic table, molecular closure, and the first cell to the connectome, the conscious episode, and the layer that states the constitution which produced it — layers 7–10 are life, layers 14–17 are mind. Qi-EP is the selector at every one of those layers: which reaction path in the chemistry layers, which flux architecture in the metabolic ones, which internal model in the cognitive ones. The capstone’s dynamics block deploys the three-gauge theorem in exactly the form proved here — the thermodynamic gauge recovering the Gibbs free energy in the $J_{Qi} = 0$ limit, the variational gauge recovering the Friston form at the declared boundary, the action gauge recovering the GG-6 field action in the detector limit — so the chain of custody runs unbroken from Eq. (32.2) to the account of why there is anyone to read it. What this Part supplies to that chain is the load-bearing middle: the selector, derived once, with its regimes attached, valid at every floor because no floor’s special furniture was used in proving it.

35 Three sciences, one principle

It is worth pausing on how odd the situation was before this Part.

Three mature sciences each had a variational principle at the bottom of it. Each principle worked. Each was arrived at independently, by people solving different problems in different centuries with different mathematics. And each field, if pressed, would have described its principle as fundamental — not as a special case of anything. A cell is simultaneously a mechanical object following a path, a thermodynamic object maintaining fluxes, and an inferential object correcting predictions; if the three principles were really three, the cell would have been receiving three sets of instructions, and there was never any account of why they did not conflict.

They did not conflict because they were one instruction. That is the content of this Part, and the form of the argument matters: the three are not shown to be analogous, or compatible, or reducible in some limit. One functional is projected from the bulk onto the boundary, and each of the three principles is recovered by asking for stationarity with a different quantity held primary. Different questions of the same object, returning different-looking answers, in the way that the same landscape gives a different-looking answer to “which way is downhill” depending on whether you are asking about a ball, a river, or a population.

It is worth being exact about what kind of unification this is, because the word is used loosely and usually means less than it sounds. Nothing here has been merged. The three principles keep their own variables, their own units and their own laboratories, and no physicist computing a trajectory will now find herself computing entropy production instead. What has changed is that they have acquired a common author. One quantity is extremized, on one surface, and the three principles are the three answers that quantity gives to three different questions. That is a weaker relation than identity and a far stronger one than analogy. Analogies are free — one can be drawn between any two things by anyone willing to squint. This costs something, because it forbids outcomes. If the three principles were really independent, a system could satisfy one and violate another, and somewhere in nature something would. The claim here is that nothing does.

The claim can now be stated at full strength, because every word of it has been earned above and each clause carries its theorem. Within this framework the principle of least time, maximum entropy production, and the free energy principle are *theorems*: each is the reading of one projected functional in its own gauge (Theorems 27.1, 28.1, 29.1, unified in Theorem 32.1), each is exact precisely where its regime theorem says, and each fails, by theorem, where the regime ends (Theorem 30.1). The derivation used nothing biological — a packet, a ledger, and a projection — so the theorems bind every open system in the SeOS class; a living organism and a brain are members by the persistence theorem of Part VII, and the companion capstone carries the demonstration up the eighteen layers from the gauge group to the sentence you are reading (Arisaka, 2026, A11-Origin). Attaching the regime to the theorem in the same breath is not a hedge; it is what makes the claim stronger than the slogans it replaces, because a theorem that knows its boundary is falsifiable and a slogan is not.

Each of the three sciences gains something specific, and the gains are worth naming separately because they are of different kinds. *Nonequilibrium thermodynamics* gains a grammar: extremum statements about constitutive relations (flux side) and extremum statements about state selection (force side) are different kinds of statement, conjugate at one point, and the field’s oldest quarrel dissolves under the distinction — with the functional now fixed upstream, so that “which entropy production” is no longer the user’s choice. *The biology of energy* gains a currency: if information carries the derived price of Qi-QUA, then bioenergetics and information processing are one subject with one ledger, and questions that sound philosophical — how much of a cell’s ATP turnover services its memory and prediction — become budgetary, and attemptable. *The sciences of mind* gain a boundary: the inference reading survives here in exactly the form its best critics said it should — a conditional description, holding where a checkable rank condition and sector condition hold, silent elsewhere. That is a demotion from “all persisting matter infers” and a promotion from “useful modeling scheme”: inference becomes a physical regime a system can be in or out of, and which regime a given preparation occupies is an experimental question.

And symmetry requires the converse list, because a Part whose subject is three literatures’ overclaims must not end by adding a fourth. It is *not* concluded that maximum entropy production is true — the derivation says it is true in a bounded regime and false outside it, and the boundedness is the result. It is *not* concluded that every organism computes probabilities — the inference reading holds where the rank condition holds and is silent elsewhere. And it is *not* concluded that the framework behind the derivation is established: the projected functional is only as trustworthy as

the six-dimensional construction, the anomaly proof, and the constitutional charter upstream of it, whose own open items are recorded in their home papers and inherited here at their stated grades (Table 26). The three principles were promoted to laws without derivation, and their fields paid for the promotion in decades of misapplication and empty quarrel; this Part's discipline is to hand every claim its passport at the border.

There is a cautionary lesson in this Part as well as a result. Two generations of argument about whether driven systems maximize or minimize entropy production dissolve here into a statement about which variables were being held fixed by whom. The disputants were not disagreeing about nature. They were disagreeing about a convention, and the disagreement lasted because there was no framework in which the convention was visible as a convention. That happens more often in science than anybody likes to admit, and it is worth knowing what it looks like from outside.

36 Concluding remarks: one landscape, three shadows

The whole Part retold, once, with no equations. Begin with the question that started it: *which way does an open system go?* Equilibrium physics has one magnificent answer for sealed boxes — downhill in free energy — and no general answer for anything alive enough to be interesting. Three sciences filled the gap with three rules: settle into the laziest steady state available (true, provably, when the pushing is gentle); degrade the driving gradient as fast as the constraints allow (striking on three planets, derived never); behave so as to make what arrives unsurprising (running half of computational neuroscience on mathematics nobody disputes, under a foundational claim that did not survive inspection). And a living cell is all three subjects at once, with no account anywhere of why it is not torn apart by three different instructions.

The answer of this Part is that there is one landscape of cost, projected down from the bulk, and the three rules are its three shadows. Ask the landscape which *path* and it answers with least time; ask it which *flows* and it answers with the flux-force law, whose two sides are the maximum and minimum principles that were never rivals; ask it which *guess* and it answers with the error-correcting rule of the brain sciences, its vague parts replaced by checkable conditions. One object, three questions, three famous answers — and, the part a skeptical reader should weigh most heavily, the failures were derived too. The circulating currents that ruin the minimum principle in a driven circuit, the convection cells that embarrassed it historically, the special couplings the brain-science principle silently needs: all of them sit exactly where the reduction says they must, in the two sectors of the flow that circulate and relabel rather than descend. A theory that predicts only successes explains nothing; the arithmetic of this one required its own limits.

What would end this Part is not a number but a system: something whose path, whose fluxes, and whose inference genuinely could not be organized by one budget — a case where the three descriptions demonstrably pull apart. The claim is that no such thing exists among systems of this kind. That is a strong claim, and it is stated in the form that lets somebody go looking; the register of §31 names the eight places to start, each with the measurement that would bite.

The reader who wants any thread at greater length has three doors. The histories and the complete proofs at review depth are the companion supplement (Arisaka, 2026, S1-EP). The cascade

that carries the selector from the gauge group to life and mind is the companion capstone ([Arisaka, 2026, A11-Origin](#)). And the remaining Parts of the present paper are the runtime that the selector serves: what the selected branch looks like under magnification, how it remembers, and what it costs to stay one.

One further thing has changed, and it sets up what follows. Before this Part the system had a budget; it now has a chosen branch — a definite path, a definite set of maintained currents, a definite model under correction. A chosen branch is an object, and objects have properties one can go and measure. The most exposed of those properties is what the branch looks like when the scale of the looking is changed, because that is the one question an instrument can ask without knowing anything about where the branch came from. The next Part asks it. It is the point at which this framework stops giving an account of things already known and starts saying things a laboratory is in a position to refuse.

Part V

Qi-CON

Conformality, Lévy statistics, fractality, and the two-knee $1/f$ band

This Part is about what a system of this kind looks like when you change the magnification.

The phenomenon at issue is one of the most reported and least explanatory in science. Measure almost any complicated thing over time — a resistor, a heartbeat, a river, a stock index, a neuron — and the fluctuations have more power at low frequencies than at high, in a particular way, over a wide range. It is called $1/f$ noise, it has been found nearly everywhere since the 1920s, and it has generated a large literature of mechanisms because so many different mechanisms can produce it. That ubiquity is exactly the problem. A prediction that a system will show an unbounded power law is nearly unfalsifiable, since almost any story delivers one.

So this Part makes a narrower claim, which is the reason it is worth making. The power-law band is *finite*, and the framework names both of its ends. Below the lower turn the spectrum flattens, and the frequency where that happens is set by the slowest process the system still keeps — the lifetime of the state it has settled into. Above the upper turn the spectrum steepens, and that frequency is set by the fastest process it still resolves. Two turns, not one; a band, not a slope. And the framework says in advance what each end should equal, which is what makes the whole thing something a measurement could contradict.

Two more things fall out of the same construction and are worth watching for, because they are one claim rather than three pieces of evidence. The jumps such a system makes should have heavy tails rather than bell-shaped ones, and the set it explores should be fractal rather than filling. The same single exponent controls the spectrum, the tails and the dimension, so a measurement that fixes it in one place fixes it in the others — and a system that gave three mutually inconsistent answers would have refuted the Part.

Two cautions before starting. Nothing here claims exact scale invariance at all scales; real systems are finite, instrument-bounded and resource-limited, and the statement is about the window between the two turns. And where the Part places known phenomena from other fields, it is placing them, not predicting them: a retrodiction of $1/f$ noise is not dressed up as a confirmation anywhere in what follows.

After budget and selector come diagnostics. Once Qi-QUA has fixed admissible current book-keeping and Qi-EP has fixed the selected history, one can finally ask what observable multiscale signatures that selected runtime leaves behind. Part V therefore derives the conformal, Lévy, fractal, and $1/f$ -type diagnostics of the selected branch.

36.1 Historical background (with citations)

Antecedents: Mandelbrot (1967, 1982), Wilson–Kogut (1974), Kadanoff (1966), Lévy (1925), Bak–Tang–Wiesenfeld (1987), Voss–Clarke (1976) (Mandelbrot, 1982, Wilson, 1971, Wilson and Kogut, 1974); Qi-CON subsumes these as gauges or limiting cases of Theorem T-CON.1 $\cup$ T-CON.2 (fractality, Lévy statistics, and the two-knee $1/f$ spectrum are joint signatures of the homogeneous boundary kernel; see A1 Sec. 6 and A4 Sec. 5 for full history).

Dual-layer architecture and the compatibility relation $\mathfrak{C}_{\text{CON}}$. Qi-CON is, in the L16 language, a dual-layer principal pillar. Its *algebraic layer* (T-CON.1) is the operational conformality theorem (Theorem 39.1): the boundary projection of the conformal isometry of the 6D bulk produces homogeneous boundary kernels with power-law correlators, realizing at the boundary the Lévy-stable distributions of Lévy (1937), the fractal geometry of Mandelbrot (1982), and the conformal field theory of Belavin, Polyakov & Zamolodchikov (1984). Its *operational/physical layer* (T-CON.2) is the two-knee spectral theorem (Theorem 42.1): the predicted $1/f$ spectrum has exactly two scale-breaking knees — the co-occurrence of two knees bracketing a finite power-law band, rather than a single unbounded slope, is the structural signature fixed by the homogeneous boundary kernel; the knee positions are the system’s slowest and fastest retained rates, and the band slope is tied to the conformal weight. This two-knee structure is the central falsifiable empirical prediction of Qi-CON. The two layers are tied together by the named compatibility relation

$$\boxed{\mathfrak{C}_{\text{CON}} \equiv \text{T-CON.1} \implies \text{T-CON.2} \quad (\text{via the scale profile and OGN-5-determined boundary kernel widths}).} \quad (36.1)$$

The bulk conformal isometry forces homogeneous boundary kernels with power-law decay; the two boundary knee positions are then set by the specific bulk scale profile and OGN-5 integers, giving a sharp falsifiable prediction.

37 Constitutional statement of the Qi-CON principle

Before any detailed calculation, the central claim of Qi-CON should be stated in a precise operational form.

Definition 37.1 (Qi-CON). Let S be an admissible open subsystem on the physical boundary state manifold $(\mathcal{X}_{\text{phys}}, g)$, and let γ_* be a Qi-EP-selected branch of its projected dynamics. We say that S is in a Qi-CON window if there exist two cutoff scales

$$\ell_{\text{micro}} \ll \ell \ll L_{\text{att}} \quad \Longleftrightarrow \quad f_{\text{IR}} \ll f \ll f_{\text{UV}}, \quad (37.1)$$

and an exponent $0 < \alpha \leq 2$ such that the effective generator of the selected branch satisfies

$$\mathcal{L}_{\text{eff}} = -D_{\alpha} (-\Delta_g)^{\alpha/2} + \mathcal{L}_{<\alpha}, \quad (37.2)$$

where $D_{\alpha} > 0$, $(-\Delta_g)^{\alpha/2}$ is the fractional Laplace–Beltrami operator on $(\mathcal{X}_{\text{phys}}, g)$, and the remainder $\mathcal{L}_{<\alpha}$ contains lower-order, cutoff-sensitive, or weakly relevant deformations. In this window the




Figure 11: The claim that makes this pillar falsifiable. Systems of this kind are known to show $1/f$ noise, and that observation on its own is worth very little, because an unbounded power law is compatible with a great many mechanisms and so eliminates almost none of them. What is predicted here is a *bounded* band: the spectrum is flat below the lower knee, follows a power law between the knees, and steepens to f^{-2} above the upper one. The two knees are not fitted. The lower sits at the slowest rate the system still keeps — the lifetime of the state it settled into — and the upper at the fastest it still resolves, so a measurement can disagree with either independently. The exponent in the band, the heaviness of the jump distribution, and the dimension of the set the system explores are all controlled by one number, so those are one claim rather than three pieces of evidence. Where the paper places known observations from other fields against this shape, they are placements: a retrodiction of $1/f$ noise is nowhere treated as a confirmation.

leading dynamics are asymptotically homogeneous under scale transformations and generate power-law correlators, heavy-tailed coarse statistics, fractal support, and $1/f$ -type spectra.

Three clarifications are essential.

First, Qi-CON does not claim exact scale invariance at all scales. Natural systems are finite, instrument-bounded, resource-limited, and often branch-selected by GSSB. Therefore the relevant statement is *intermediate* or *operational* conformality: there exists a nontrivial scale interval on which the leading generator is homogeneous to good approximation.

Second, the exponent α is not inserted by hand. It is inherited from the projected bulk-to-boundary kernel and therefore ultimately from the six-dimensional geometry together with the selected branch. The value $\alpha = 2$ reproduces Gaussian diffusion as a limiting case, but the generic Qi-CON regime of interest is $0 < \alpha < 2$, where nonlocal heavy-tailed transport survives at leading order.

Third, the phrase “conformality” is used in the operational sense appropriate to GDOS. In the strictest mathematical literature, full conformal invariance can be stronger than mere scale covariance. The present paper uses the term to denote the physically relevant combination of homogeneous generators, power-law correlators, and absence of an intrinsic scale in the active window. This is exactly the feature required for a universal diagnostic layer.

Remark 37.2. The transition from exact microscopic law to operational conformality mirrors the transition from exact bulk action to projected boundary action already encountered in Qi-EP. In both cases the decisive point is not that the boundary is simpler than the bulk, but that the observer sees only the projected homogeneous core together with finite-size and instrument-induced corrections.

Table 28: Four canonical signatures of a Qi-CON window. The rows are operational conformality, Lévy statistics, fractality, and $1/f$ noise with two knees; the columns give the mathematical form of each and its meaning inside the boundary law. They are listed together because they are not four independent findings — the first three are one exponent seen three ways, as Table 29 sets out, and the fourth is the bounded band that distinguishes this claim from an unbounded power law.

Signature	Mathematical form	Meaning inside GDOS
Operational conformality	$C(\lambda z, \lambda^\alpha t) = \lambda^{-2\Delta} C(z, t)$	No intrinsic scale inside the active window; correlations decay by power laws.
Lévy statistics	$\hat{p}(k, t) = \exp[-D_\alpha t k ^\alpha]$	Coarse increments are α -stable rather than Gaussian.
Fractality	$N(\varepsilon) \sim \varepsilon^{-D_f}$	Trajectories and selected supports occupy non-integer effective dimension.
$1/f$ with two knees	$S(f) \propto f^{-\beta}$ on $f_{\text{IR}} < f < f_{\text{UV}}$	The selected branch produces multiscale soft modes bounded by attractor and microphysical cutoffs.

38 From the 6D bulk to homogeneous boundary kernels

The derivation of Qi-CON begins not on the boundary but in the six-dimensional bulk already established in GG-A2-GG6 (Arisaka, 2026, A2-GG6). The extra radial direction is not an ornamental coordinate. It is the geometric carrier of scale. Writing

$$r = \ln\left(\frac{E_{\text{GUT}}}{\mu}\right), \quad (38.1)$$

the ultraviolet boundary corresponds to $r = 0$ and lower operational scales correspond to larger r . In other words, the bulk geometry stores the renormalization structure that later appears on the boundary as multiscale dynamics (Arisaka, 2026, A2-GG6; Arisaka, 2026, A5-GDOS; Wilson, 1971; Wilson and Kogut, 1974; Maldacena, 1998).

To expose the scaling law cleanly, it is convenient to replace r by a positive scale coordinate

$$\zeta = e^{-r/L}, \quad (38.2)$$

where L is the radial length scale. Large ζ corresponds to coarse boundary scales, whereas small ζ corresponds to ultraviolet scales. Consider a projected Qi-neutral observable on a local physical state patch $U \subset \mathcal{X}_{\text{phys}}$ of dimension

$$d_{\mathcal{X}} = \dim U. \quad (38.3)$$

Let $\ell = d_g(z, z')$ denote the geodesic separation of two nearby boundary points $z, z' \in U$. For a dominant projected bulk mode with effective scaling weight Δ , the bulk-to-boundary propagator has the generic scale-profile form

$$K_{\Delta}(\zeta; z, z') = \left(\frac{\zeta}{\zeta^2 + \ell^2}\right)^{\Delta}, \quad (38.4)$$

up to slowly varying geometric prefactors (Maldacena, 1998). The essential feature is homogeneity: the kernel depends on z, z' only through the dimensionless ratio ℓ/ζ .

The OGGEp projection integrates over the hidden scale coordinate and the unresolved sector. A jump *rate* between two boundary points is an amplitude squared: the projected mode reaches each endpoint of the jump through one factor of the bulk-to-boundary propagator, so the induced boundary jump kernel carries K_{Δ} twice,

$$W_{\Delta}(z, z') = \int_{\zeta_{\min}}^{\zeta_{\max}} \omega(\zeta) [K_{\Delta}(\zeta; z, z')]^2 d\zeta, \quad (38.5)$$

where $\omega(\zeta)$ is the effective radial weight after KK summation, scale profile Jacobians, and instrument projection. In the active scaling window one may write

$$\omega(\zeta) = \zeta^{\sigma-1} w(\zeta), \quad w(\zeta) \approx w_0, \quad (38.6)$$

with slowly varying w_0 . Substituting Eq. (38.4) into Eq. (38.5) gives

$$W_\Delta(z, z') \approx w_0 \int_{\zeta_{\min}}^{\zeta_{\max}} \frac{\zeta^{\sigma+2\Delta-1}}{(\zeta^2 + \ell^2)^{2\Delta}} d\zeta. \quad (38.7)$$

Now perform the scale-separating substitution

$$\zeta = \ell u, \quad d\zeta = \ell du. \quad (38.8)$$

One obtains

$$W_\Delta(z, z') \approx w_0 \ell^{\sigma-2\Delta} \int_{\zeta_{\min}/\ell}^{\zeta_{\max}/\ell} \frac{u^{\sigma+2\Delta-1}}{(1+u^2)^{2\Delta}} du. \quad (38.9)$$

Inside the operational conformal window,

$$\zeta_{\min} \ll \ell \ll \zeta_{\max}, \quad (38.10)$$

the integral in Eq. (38.9) saturates to the finite dimensionless constant

$$C_{\Delta,\sigma} = \int_0^\infty \frac{u^{\sigma+2\Delta-1}}{(1+u^2)^{2\Delta}} du = \frac{1}{2} B\left(\Delta + \frac{\sigma}{2}, \Delta - \frac{\sigma}{2}\right), \quad |\sigma| < 2\Delta, \quad (38.11)$$

the convergence condition and the admissibility bounds below being the same fact seen twice, and therefore

$$W_\Delta(z, z') \propto \ell^{-(2\Delta-\sigma)}. \quad (38.12)$$

This is the decisive scaling step: after the hidden radial coordinate has been integrated out, the boundary inherits a power-law kernel with no intrinsic scale of its own. The full algebra is collected in Appendix G.1.

To compare Eq. (38.12) with standard nonlocal generators, write

$$2\Delta - \sigma = d_{\mathcal{X}} + \alpha, \quad 0 < \alpha \leq 2. \quad (38.13)$$

The admissibility window $0 < \alpha \leq 2$ is forced as follows. The upper bound $\alpha = 2$ is the Gaussian (heat-kernel) limit: there the fractional Laplacian $(-\Delta_g)^{\alpha/2}$ reduces to the ordinary Laplacian, the boundary kernel becomes a Brownian semigroup, and the Lévy parameter degenerates to a Gaussian (Lévy, 1937). Values $\alpha > 2$ fail to define a non-negative generator on physical states: the corresponding kernel $d_g(z, z')^{-(d_{\mathcal{X}}+\alpha)}$ is more singular at $z' = z$ than the principal-value integral can absorb, and the resulting operator is no longer dissipative (standard result on Riesz potentials). The lower bound $\alpha > 0$ excludes the degenerate case in which the kernel ceases to be locally integrable and the operator collapses to a constant. In terms of the bulk parameters, the admissibility window translates to $\Delta > (d_{\mathcal{X}} + \sigma)/2$ (lower bound) and $2\Delta - \sigma \leq d_{\mathcal{X}} + 2$ (upper bound). Both bounds are

routinely respected by the operationally selected branches discussed in the rest of the section. Then

$$W_\Delta(z, z') \propto \frac{1}{d_g(z, z')^{d_\chi + \alpha}}, \quad (38.14)$$

which is exactly the homogeneity required for a fractional jump operator on U .

Proposition 38.1 (Homogeneous kernel induced by OGGE). *Assume that a Qi-EP-selected branch admits an operational window satisfying Eq. (38.10) and that the projected radial weight is slowly varying there. Then the OGGE-induced nonlocal kernel on the corresponding physical boundary patch is asymptotically homogeneous with degree $-(d_\chi + \alpha)$. Consequently the leading nonlocal part of the effective generator is*

$$(\mathcal{J}_\alpha f)(z) = C_{d_\chi, \alpha} \text{PV} \int_U \frac{f(z') - f(z)}{d_g(z, z')^{d_\chi + \alpha}} d\mu_g(z'), \quad (38.15)$$

which, in a local flat patch, reduces to

$$\mathcal{J}_\alpha = -D_\alpha (-\Delta_g)^{\alpha/2}. \quad (38.16)$$

Proof. Equation (38.9) already shows that the projected kernel is homogeneous with degree $-(2\Delta - \sigma)$. Identifying that degree with $d_\chi + \alpha$ gives Eq. (38.14). The principal-value integral (38.15) is the unique translation-covariant jump operator with that homogeneity on a local Euclidean patch, and the curved-space generalization is the fractional Laplace–Beltrami operator to leading order in the local chart. Thus the nonlocal sector of the boundary generator has the form (38.16). ■

The proposition explains the geometric origin of the exponent α . It is not introduced phenomenologically by fitting a heavy-tailed random walk to data. It is the boundary image of the hidden-scale integral. Once the selected branch fixes the dominant mode Δ and the effective measure exponent σ , the exponent α is fixed.

39 The leading Qi-CON generator and operational conformality

On the selected branch γ_* , the most general projected boundary equation combines drift, local diffusion, and the nonlocal term derived above:

$$\partial_t \rho = -\nabla_a (u^a \rho) + \nabla_a \nabla_b (D^{ab} \rho) - D_\alpha (-\Delta_g)^{\alpha/2} \rho + \mathcal{S}_{\text{ext}}[\rho]. \quad (39.1)$$

Here u^a collects branch-dependent drift or transport, D^{ab} is the local diffusion tensor, and $\mathcal{S}_{\text{ext}}$ denotes sources, sinks, or instrument-specific deformations. Equation (39.1) is exact enough for interpretation but still too detailed for a clean scaling analysis.

The Qi-CON window is precisely the regime in which the homogeneous nonlocal sector dominates. This occurs when the observational scale is much larger than the microscopic local-diffusion scale but much smaller than the attractor-confinement scale. In that regime one may write

$$\partial_t \rho = -D_\alpha (-\Delta_g)^{\alpha/2} \rho + \text{sub-leading terms}. \quad (39.2)$$

The scale covariance of the leading term is immediate. Let

$$(\mathcal{S}_\lambda f)(z) = f(\lambda z), \quad \lambda > 0. \quad (39.3)$$

Then a direct Fourier or coordinate calculation yields

$$\mathcal{S}_\lambda^{-1}(-\Delta_g)^{\alpha/2}\mathcal{S}_\lambda = \lambda^\alpha(-\Delta_g)^{\alpha/2}, \quad (39.4)$$

modulo curvature corrections that are suppressed on the local scaling patch. Therefore Eq. (39.2) is invariant under the simultaneous transformation

$$z \mapsto \lambda z, \quad t \mapsto \lambda^\alpha t. \quad (39.5)$$

This is the core of operational conformality: the selected open-system dynamics possess no preferred length or time scale inside the active interval.

A point-source initial condition makes the scaling explicit. In a local flat patch, Fourier transformation of Eq. (39.2) gives

$$\partial_t \widehat{\rho}(k, t) = -D_\alpha |k|^\alpha \widehat{\rho}(k, t), \quad (39.6)$$

so that

$$\widehat{\rho}(k, t) = e^{-D_\alpha t |k|^\alpha} \widehat{\rho}(k, 0). \quad (39.7)$$

For $\rho(z, 0) = \delta(z)$, inverse transformation yields the scaling kernel

$$p_\alpha(t, z) = t^{-d_\chi/\alpha} \Phi_\alpha\left(\frac{z}{t^{1/\alpha}}\right), \quad (39.8)$$

where Φ_α is a dimensionless profile determined by α and the local metric. Equation (39.8) is the precise semigroup form of Qi-CON scaling.

Theorem 39.1 (Operational conformality theorem (T-Qi4-11.4)). *Suppose that a Qi-EP-selected branch admits the Qi-CON leading equation (39.2) on a local physical state patch of dimension d_χ . Then the fundamental solution satisfies the homogeneity law*

$$p_\alpha(\lambda^\alpha t, \lambda z) = \lambda^{-d_\chi} p_\alpha(t, z). \quad (39.9)$$

Consequently, for any scale-covariant Qi-neutral observable O with effective boundary dimension Δ_O , the two-point function in the active window obeys

$$\langle O(\lambda z, \lambda^\alpha t) O(0, 0) \rangle = \lambda^{-2\Delta_O} \langle O(z, t) O(0, 0) \rangle, \quad (39.10)$$

and in particular at equal time

$$\langle O(z) O(0) \rangle \propto |z|^{-2\Delta_O}. \quad (39.11)$$

Proof. Equation (39.9) follows directly from Eq. (39.8) by rescaling t and z . Equation (39.10) then follows from dimensional covariance of the semigroup acting on observables with effective weight Δ_O . Setting $t = 0$ gives Eq. (39.11). ■

The theorem is intentionally phrased in terms of an operational window rather than a global conformal field theory. That is the natural level for GDOS. The observer does not access the full microscopic theory directly; the observer accesses the homogeneous boundary image of the selected branch over a finite but nontrivial scale range.

40 Second signature: Lévy statistics from the fractional generator

The second signature of Qi-CON is the appearance of α -stable coarse statistics. Once the leading generator has the fractional form (39.2), the characteristic function of a scalar coarse coordinate follows immediately. Let X_t denote one local coordinate projection of the boundary state. In the drift-free leading regime its density $p(x, t)$ satisfies

$$\partial_t p(x, t) = -D_\alpha (-\partial_x^2)^{\alpha/2} p(x, t). \quad (40.1)$$

Fourier transformation gives

$$\partial_t \hat{p}(k, t) = -D_\alpha |k|^\alpha \hat{p}(k, t), \quad (40.2)$$

with solution

$$\hat{p}(k, t) = \exp[-D_\alpha t |k|^\alpha] \hat{p}(k, 0). \quad (40.3)$$

For a localized initial packet $\hat{p}(k, 0) = 1$, the characteristic function becomes

$$\phi_{X_t}(k) = \mathbb{E}[e^{ikX_t}] = \exp[-D_\alpha t |k|^\alpha]. \quad (40.4)$$

This is the defining characteristic function of a symmetric α -stable law. The full derivation is written out in Appendix G.2 and agrees with the standard anomalous-diffusion literature (Metzler and Klafter, 2000, van Kampen, 2007).

Several consequences follow at once.

First, the process is self-similar:

$$X_t \stackrel{d}{=} t^{1/\alpha} X_1. \quad (40.5)$$

Second, for $0 < \alpha < 2$ the tails are heavy:

$$p(x, t) \sim \frac{D_\alpha t \Gamma(1 + \alpha) \sin(\pi\alpha/2)}{\pi |x|^{1+\alpha}} \quad (|x| \rightarrow \infty), \quad (40.6)$$

so the coarse increments are Lévy rather than Gaussian. Third, the Gaussian case appears only as the codimension-one limit $\alpha = 2$, where Eq. (40.4) becomes $e^{-D_2 t k^2}$.

Corollary 40.1 (Why heavy tails are generic inside Qi-CON). *Within an operational Qi-CON window, the dominant boundary increments are generically heavy-tailed whenever $0 < \alpha < 2$.*

Exact Gaussianity requires $\alpha = 2$, which corresponds to the disappearance of the leading nonlocal contribution.

Proof. Equation (40.4) is α -stable for all $0 < \alpha \leq 2$. For $0 < \alpha < 2$, the inverse Fourier transform has the asymptotic power law (40.6); for $\alpha = 2$ the law becomes Gaussian. Hence heavy tails are generic except on the Gaussian boundary. ■

Physically, this means that selected open-system trajectories retain rare but consequential excursions across scales. The system does not explore its branch by small equal-sized Brownian steps alone. It mixes local relaxation with occasional larger state-space moves. In an exact mathematical model these are Lévy flights; in realistic finite systems they are often truncated or tempered by the same cutoffs that define the operational window. The present framework therefore predicts *approximately* α -stable increments over the active window, not an infinite-scale law.

41 Third signature: fractality of trajectories and selected supports

The third Qi-CON signature is fractality. This statement should not be treated as a poetic metaphor. It follows from the same self-similarity exponent that controlled the characteristic function.

Let R_T denote the range of a coarse trajectory on a finite observation interval:

$$R_T = \{X_t \mid 0 \leq t \leq T\} \subset U. \quad (41.1)$$

To estimate its box-counting dimension, cover R_T with boxes of side ε . Because the process is self-similar with exponent $1/\alpha$, a time interval δt typically produces a spatial excursion of size

$$\delta x \sim (\delta t)^{1/\alpha}. \quad (41.2)$$

To resolve the path at spatial resolution ε , choose $\delta x \sim \varepsilon$, which implies

$$\delta t \sim \varepsilon^\alpha. \quad (41.3)$$

An observation window of total duration T therefore contains approximately

$$N_T(\varepsilon) \sim \frac{T}{\delta t} \sim T \varepsilon^{-\alpha} \quad (41.4)$$

independent ε -scale segments. This estimate cannot exceed the number of ε -boxes available in the ambient physical patch, which scales as $\varepsilon^{-d_{\mathcal{X}}}$. Therefore the effective box-counting dimension of the range is

$$D_f = \min(d_{\mathcal{X}}, \alpha). \quad (41.5)$$

The derivation is elementary but important enough that it is recorded separately in Appendix G.3; it is also consistent with the standard fractal viewpoint (Mandelbrot, 1982).

Equation (41.5) immediately explains why Qi-CON belongs in a theory of selected open systems rather than in a theory of unrestricted random walks. The ambient manifold dimension $d_{\mathcal{X}}$ sets

the maximal available space, while the exponent α sets the intrinsic scaling of the selected path. The realized fractal support is the lesser of the two. When GSSB further constrains the branch to a submanifold or attractor sheet, the effective support dimension can only decrease. Thus the fractality seen by the observer is not merely a property of the transport kernel; it is also a property of branch selection and attractor confinement.

Proposition 41.1 (Fractal support of a selected Qi-CON branch). *Let a Qi-EP-selected branch possess a Qi-CON window with exponent $0 < \alpha \leq 2$ on an ambient physical patch of dimension $d_{\mathcal{X}}$. Then the coarse trajectory range satisfies the box-counting estimate*

$$N(\varepsilon) \sim \varepsilon^{-D_f}, \quad D_f = \min(d_{\mathcal{X}}, \alpha), \quad (41.6)$$

over $\ell_{\text{micro}} < \varepsilon < L_{\text{att}}$. If GSSB restricts the branch to an attractor support $A \subset U$, then the observable support dimension satisfies $D_A \leq D_f$.

Proof. The range estimate follows from Eqs. (41.2)–(41.5). The second statement follows because any attractor support A visited by the selected branch is a subset of the ambient range and therefore cannot have a larger box-counting dimension than the range itself. ■

A useful way to phrase the result is that Qi-CON predicts *fractal admissibility*. The selected branch need not fill the ambient state space densely; indeed it usually does not. It occupies a support whose effective dimension is determined jointly by the scale exponent and by the symmetry-breaking constraints that shaped the branch.

Table 29: How a single exponent organizes the first three Qi-CON signatures. The semigroup kernel, the characteristic function and the fractal support occupy the three rows, with the scaling law and the interpretation alongside. The table is the reason the signatures are not independent evidence: the same α appears in all three, so a measurement that fixes it in one row fixes it in the others, and a framework claiming all three is making one claim rather than three.

Object	Scaling law	Interpretation
Semigroup kernel	$p_{\alpha}(t, z) = t^{-d_{\mathcal{X}}/\alpha} \Phi_{\alpha}(z/t^{1/\alpha})$	Operational conformality and absence of intrinsic scale in the active window.
Characteristic function	$\phi_{X_t}(k) = e^{-D_{\alpha} t k ^{\alpha}}$	α -stable Lévy statistics for coarse increments.
Fractal support	$D_f = \min(d_{\mathcal{X}}, \alpha)$	Non-integer effective dimension of selected trajectory ranges or supports.

42 Fourth signature: $1/f$ spectra with two knees

The fourth Qi-CON signature is the most directly measurable in many applications: a broad spectral band of the form $1/f^{\beta}$, bounded by two characteristic knees. The derivation proceeds most transparently from the spectral representation of the selected branch generator.

Let $X(t)$ be a centered observable measured on the selected branch. Write its autocorrelation function as a mode expansion over the nonnegative spectrum of the reduced generator:

$$C_X(\tau) = \langle X(t+\tau)X(t) \rangle = \int_{\lambda_{\text{IR}}}^{\lambda_{\text{UV}}} A_X(\lambda) e^{-\lambda|\tau|} \rho(\lambda) d\lambda. \quad (42.1)$$

Here $\rho(\lambda)$ is the density of relaxation modes, $A_X(\lambda)$ is the observable-dependent spectral weight, λ_{IR} is the slowest relevant mode, and λ_{UV} is the fastest mode retained by the operational description. Inside the Qi-CON window, the same homogeneity that produced the fractional operator implies a power-law soft-mode density:

$$A_X(\lambda)\rho(\lambda) \approx A_0 \lambda^{\gamma-1}, \quad 0 \leq \gamma < 1, \quad (42.2)$$

for $\lambda_{\text{IR}} < \lambda < \lambda_{\text{UV}}$. The exponent γ is not phenomenological: it is fixed by Weyl asymptotics of the fractional Laplace–Beltrami operator on a $d_{\mathcal{X}}$ -dimensional manifold. Indeed, the eigenvalues λ_n of the standard Laplacian on a compact $d_{\mathcal{X}}$ -manifold satisfy $\lambda_n \sim n^{2/d_{\mathcal{X}}}$, so that the eigenvalue counting function $N(\lambda) \sim \lambda^{d_{\mathcal{X}}/2}$ and the density $\rho_{\Delta_g}(\lambda) \sim \lambda^{d_{\mathcal{X}}/2-1}$ (Weyl’s law). The fractional operator $(-\Delta_g)^{\alpha/2}$ acts on the same eigenspaces with eigenvalues $\lambda_n^{\alpha/2}$; changing variables yields the fractional-operator density

$$\rho_{\alpha}(\lambda) \sim \lambda^{d_{\mathcal{X}}/\alpha-1}. \quad (42.3)$$

If the observable spectral weight scales as $A_X(\lambda) \sim \lambda^{\beta-d_{\mathcal{X}}/\alpha}$ on the soft modes (the natural scaling for any observable that couples through the gradient of the fractional operator), then $A_X(\lambda)\rho_{\alpha}(\lambda) \sim \lambda^{\beta-1}$, identifying $\gamma = \beta$ in Eq. (42.2). The nearly critical case $\gamma \approx 0$ corresponds to $\beta \approx 0$, i.e. an observable that couples uniformly across the soft-mode spectrum — the universal generic case for memory-bearing systems, and the one that leads to true $1/f$ behavior.

The power spectral density is

$$S_X(\omega) = 2 \int_0^{\infty} C_X(\tau) \cos(\omega\tau) d\tau. \quad (42.4)$$

Substituting Eq. (42.1) and evaluating the Laplace transform of each exponential mode yields

$$S_X(\omega) = 2 \int_{\lambda_{\text{IR}}}^{\lambda_{\text{UV}}} \frac{\lambda A_X(\lambda)\rho(\lambda)}{\lambda^2 + \omega^2} d\lambda. \quad (42.5)$$

Using the power-law density (42.2), one finds

$$\begin{aligned} S_X(\omega) &\approx 2A_0 \int_{\lambda_{\text{IR}}}^{\lambda_{\text{UV}}} \frac{\lambda^{\gamma}}{\lambda^2 + \omega^2} d\lambda \\ &= 2A_0 \omega^{\gamma-1} \int_{\lambda_{\text{IR}}/\omega}^{\lambda_{\text{UV}}/\omega} \frac{u^{\gamma}}{1 + u^2} du, \quad u = \frac{\lambda}{\omega}. \end{aligned} \quad (42.6)$$

Now the three frequency regimes separate cleanly.

Intermediate band: the $1/f^\beta$ law. If

$$\lambda_{\text{IR}} \ll \omega \ll \lambda_{\text{UV}}, \quad (42.7)$$

the integration limits in Eq. (42.6) may be extended to 0 and ∞ , giving a constant dimensionless integral:

$$\int_0^\infty \frac{u^\gamma}{1+u^2} du = \frac{\pi}{2} \sec\left(\frac{\pi\gamma}{2}\right), \quad 0 \leq \gamma < 1. \quad (42.8)$$

Therefore

$$S_X(\omega) \propto \omega^{\gamma-1} = \frac{1}{\omega^\beta}, \quad \beta = 1 - \gamma. \quad (42.9)$$

The nearly critical case $\gamma \approx 0$ gives $\beta \approx 1$, i.e. genuine $1/f$ behavior.

Low-frequency side: the infrared knee. If $\omega \ll \lambda_{\text{IR}}$, then $\lambda^2 + \omega^2 \approx \lambda^2$, so

$$S_X(\omega) \approx 2A_0 \int_{\lambda_{\text{IR}}}^{\lambda_{\text{UV}}} \lambda^{\gamma-2} d\lambda = \text{constant}. \quad (42.10)$$

Thus the spectrum flattens at low frequency. In ordinary language, there is a largest effective relaxation time beyond which the selected branch behaves as a finite system rather than as a scale-free one.

High-frequency side: the ultraviolet knee. If $\omega \gg \lambda_{\text{UV}}$, then $\lambda^2 + \omega^2 \approx \omega^2$, so

$$S_X(\omega) \approx \frac{2A_0}{\omega^2} \int_{\lambda_{\text{IR}}}^{\lambda_{\text{UV}}} \lambda^\gamma d\lambda \propto \omega^{-2}. \quad (42.11)$$

Thus the spectrum steepens above the active window and leaves the scale-free regime.

These three regimes may be written compactly as

$$S_X(f) \approx \begin{cases} S_0, & f < f_{\text{IR}}, \\ A f^{-\beta}, & f_{\text{IR}} < f < f_{\text{UV}}, \\ B f^{-2}, & f > f_{\text{UV}}, \end{cases} \quad (42.12)$$

with

$$f_{\text{IR}} = \frac{\lambda_{\text{IR}}}{2\pi}, \quad f_{\text{UV}} = \frac{\lambda_{\text{UV}}}{2\pi}. \quad (42.13)$$

The detailed algebra is given in Appendix G.4. The key point is that the two knees are not arbitrary fitting decorations. They are forced by the existence of both an attractor-scale cutoff and a microphysical cutoff.

Theorem 42.1 (Two-knee spectral theorem (T-Qi4-11.7)). *Suppose a Qi-EP-selected branch possesses a Qi-CON window with soft-mode density (42.2). Then the power spectral density of any observable with nonvanishing overlap on that soft sector obeys Eq. (42.12). The infrared knee is fixed by the*

longest branch or attractor relaxation time, and the ultraviolet knee is fixed by the shortest retained microscopic time.

Proof. Equations (42.5)–(42.11) already provide the three asymptotic regimes. The infrared plateau is controlled by the lower modal cutoff λ_{IR} , while the ultraviolet roll-off is controlled by the upper cutoff λ_{UV} . Translating to ordinary frequency gives Eq. (42.13). ■

This theorem is central because it converts a broad and often qualitative empirical observation — “many open systems show $1/f$ -like spectra” — into a structural consequence of the selected branch geometry. The scale-free band is not universal by accident. It appears because the selected branch carries a dense soft-mode sector between a largest confinement scale and a smallest microphysical scale.

Table 30: Physical meaning of the two Qi-CON knees. The infrared and ultraviolet knees occupy the rows, with the mathematical origin of each and its physical interpretation alongside. The two-knee structure is what separates this prediction from the observation of $1/f$ noise, which is ubiquitous and explains little: an unbounded slope is compatible with almost anything, whereas a band bounded at both ends by scales the framework computes is falsifiable, and the second column says what sets each end.

Knee	Mathematical origin	Physical interpretation
Infrared knee f_{IR}	Lowest retained relaxation rate λ_{IR}	Largest branch or attractor timescale; beyond this the finite selected system behaves as a bounded object rather than a scale-free one.
Ultraviolet knee f_{UV}	Highest retained relaxation rate λ_{UV}	Fastest microscopic or instrument-resolved process; above this the effective coarse description breaks down and the spectrum steepens.

43 Why the Qi-CON exponents are runtime invariants

The four signatures derived above do not float independently of the master runtime: they are constrained by DGI, GSSB, GA, Qi-QUA, and Qi-EP. DGI acts on the effective generator by similarity (gauge conjugation), so the exponents α, γ, β and the knee locations are spectral invariants of the selected branch. GSSB fixes the universality class: different broken branches carry their own exponent sets $(\alpha_i, \gamma_i, \beta_i, D_{f,i}, f_{\text{IR},i}, f_{\text{UV},i})$. GA fixes the infrared knee through $f_{\text{IR}} \sim 1/\tau_{\text{GA}}$, tying it to the slowest attractor relaxation. Qi-QUA delivers spectral locking: in the ideal locked regime $J_E/Q_B = J_e = J_B = J_Q$ implies $S_{J_E}(f)/Q_B^2 = S_{J_e}(f) = S_{J_B}(f) = S_{J_Q}(f)$ across any common Qi-CON window, with realistic overhead leaving the slopes β common. Finally, Qi-EP supplies the selector that picks which spectrum is actually observed. This ordering—budget $\rightarrow$ selector $\rightarrow$ diagnostic—is why the pillars cannot be permuted. Table 31 records the one-row-per-construct summary.

Table 31: How the earlier runtime layers constrain Qi-CON. Five earlier layers run down the first column — the three cascade stages and the two preceding pillars — with the constraint each imposes and the consequence for the scale-coherence statement. The table is the pillar’s dependency surface: it inherits rather than assumes, and a reader auditing whether the third pillar is independent of the first two should read the middle column, which says exactly where it is not.

Earlier layer	Constraint on Qi-CON	Consequence
DGI	Similarity-covariant generator and gauge-invariant observables	Scaling exponents are branch invariants, not gauge artifacts.
GSSB	Branch and support selection	Different broken branches can carry different exponent sets and knees.
GA	Attractor persistence and return times	The infrared knee is fixed by the slowest attractor relaxation mode.
Qi-QUA	Common current denomination	Sector-resolved amplitudes and cross-spectral ratios are tied by the Arizaka Equation.
Qi-EP	Extremal branch selection	The observed spectrum is the spectrum of the selected branch.

44 Why Qi-CON applies to all open systems

The universality of Qi-CON now follows from a simple structural observation. Any open system with

- (i) a nontrivial separation between microscopic and macroscopic scales,
- (ii) a selected persistent branch or attractor,
- (iii) repeated throughput of energy, matter, entropy, or information,
- (iv) and observer access only to a projected coarse state

is a candidate for an operational Qi-CON window. No special reference to particles, chemistry, or brains is needed at this level. The statement is geometric and statistical.

This does *not* mean that every dataset should exhibit exact power laws over arbitrarily many decades. Finite-size effects, external forcing, insufficient scale separation, coarse instrument bandwidth, and branch switching can all shorten or obscure the active window. The point is subtler and more useful: whenever a selected open system possesses a clear intermediate scale regime, the natural generic diagnostic of that regime is the Qi-CON quartet — power-law correlations, heavy tails, fractal support, and a two-knee $1/f^\beta$ spectrum.

A domain table makes the statement concrete.

Table 32: Illustrative cross-domain readings of Qi-CON. Five domains run down the first column, with the typical coarse observable in each, the knees one would expect to find, and the expected readout. As with Table 27, these are placements rather than predictions — the paper is explicit that it does not dress a retrodiction of $1/f$ noise as a confirmation — and the value of the table is that it says in advance where the framework expects to find *two* knees rather than one.

Domain	Typical coarse observable	Likely knees	Expected Qi-CON readout
Transport or particle-flow systems	Arrival-time, channel current, branch occupancy	Device size / microscopic scattering time	Power-law transit statistics and scale-covariant throughput fluctuations.
Molecular machines	Dwell time, rotational step interval, cycle current	Whole-cycle return time / chemical substep time	Heavy-tailed dwell distributions and broad intermediate spectral band.
Cells and reaction networks	Flux, burst size, concentration time series	Homeostatic relaxation / reaction microtime	Fractal switching, broad throughput spectra, and correlated fluctuations across scales.
Brains and cognitive systems	Field potentials, avalanche size, decision-variable trajectory	Behavioral or theta-cycle time / synaptic or cellular time	$1/f$ -like spectra with two knees, heavy-tailed event statistics, and scale-free internal dynamics.
Ecological and social systems	Population change, traffic, transaction or communication counts	Institutional or seasonal time / individual-event time	Broad multiscale fluctuation windows with truncated heavy tails and knee-bounded spectra.

The table is intentionally broad because broadness is the point. Qi-CON is not a domain-specific fit function. It is a universal diagnostic grammar for selected open-system branches.

45 Why two knees are worth more than one slope

There is a way that a theory can be popular and useless at the same time, and $1/f$ noise is the standard example.

It is everywhere. It has been reported in semiconductors, in nerve cells, in the flow of the Nile, in traffic, in music, in the timing of human keystrokes. Dozens of mechanisms generate it. And that is exactly why finding it in a new place is worth very little: a prediction that something will show a power law is compatible with an enormous number of underlying stories, so confirming it eliminates almost nothing. A theory that predicted $1/f$ noise and stopped there would be collecting agreement without ever risking anything.

The move this Part makes is to give the band edges. A slope with no ends can be fitted to almost anything; a band with two ends, each fixed by a quantity the framework identifies independently, can be wrong in several distinct ways at once. It can be wrong because there is only one turn. It can be wrong because the lower turn does not sit at the lifetime of the state the system settled into. It can be wrong because the upper turn does not sit at the fastest process the system still resolves. It can be wrong because the exponent in the band, the shape of the jump distribution,

and the dimension of the set the system explores disagree about the single number that is supposed to control all three.

Underneath that is a point about how much work one number is being asked to do. This Part does not offer four predictions. It offers one, wearing four costumes. The steepness of the spectrum inside the band, the weight of the tail on the distribution of jumps, the roughness of the set the system actually wanders over, and the placement of the two turns are not four independent findings that happened to be made in the same laboratory. They descend from a single exponent. That is a dangerous position to occupy and it is occupied deliberately, because four quantities controlled by one number can be checked against each other by somebody who does not believe a word of the theory that produced them. A framework whose predictions were independent could lose one and keep the others, and would quietly do so. This one cannot. Untie any pair of these and the whole comes apart.

That is the difference between a framework that accommodates a phenomenon and one that is exposed to it. Accommodation is cheap and this Part deliberately declines it: where the domain tables place known observations from other fields, they are placements, and nothing here treats a retrodiction of $1/f$ noise as a confirmation of anything.

There is something in the bare idea of scale covariance worth pausing over before leaving this Part. To say that a system looks the same at many scales is to say that it does not carry the information of how big it is. Most things carry that information and cannot help it: a table has a length, an atom has a radius, and neither can be mistaken for a larger or smaller version of itself. What is unusual about the systems this framework is about is that between their two turns they forget their own size — and it is exactly inside that forgetful interval that they are able to do the things that make them worth studying. A structure with a preferred scale can only ever be that structure. A structure without one can be a template, and a template is what gets used again at a size nobody planned for.

A closing thought about what the two turns actually mean. They are the two ends of a system's world — the slowest thing it still is, and the fastest thing it can still tell apart. Everything between them is where that system lives, and the width of that interval is a measure of how much room it has to be complicated in. The next Part will use exactly that width to say how deep a memory such a system can build. It is an unusually direct link between something an instrument reads off a spectrum and something architectural about the thing being measured, and it is the kind of link a framework has to produce if it is going to earn the reach it claims.

Part VI

Qi-HAL

Memory architecture: MePMoS, NHT, HAL, BFFT, and the GRIP-loop

The moment a system holds on to something for longer than it takes to act and look again, it stops being merely a thing that reacts and starts being a thing with a past. This Part is about what that costs and what it forces.

The entrance condition is exactly that comparison of two times: how long a trace survives, against how long one cycle of acting and observing takes. Below it, memory cannot guide anything, because whatever was stored is gone before it could be used. Above it, the Part shows, the system's runtime falls into a definite shape — a loop that runs from memory to prediction to motion to sensing and back; a code that stores relationships between things rather than the things themselves; a lattice rather than a filing cabinet; and a ladder of nested timescales whose depth is fixed by the spectral window of the previous Part.

The order of that loop is not a modeling choice. Put sensing before action and the system would have to know an observation it has not yet caused. Put prediction after sensing and prediction becomes hindsight. Put the memory update before the new measurement arrives and the error-correction step disappears. Each rearrangement destroys something, which is why the order comes out the way it does.

Two consequences deserve to be seen before the derivations. Because what is stored is a set of differences between phases rather than absolute values, a memory written on one carrier can be read out on a completely different one — which is the precise reason a chemical system, a neural system and a digital one can hold the same content on utterly different hardware. And because the content is spread across relations rather than filed in places, a damaged substrate can still reconstruct part of what it held, which is a familiar fact about brains given a mechanism here.

Memory is never free in this account. A write is an elementary event, and the ledger of the first pillar bills it. That is the sentence that rules out the comfortable fiction of information as something symbolic floating above the physics.

One note on the strength of the claim, because the Part is careful about it and the introduction should be too. What is established is a canonical normal form — the shape a memory-bearing system's runtime takes — rather than a proof that no other shape could exist. The Part says so itself, and explains why the stronger argument would be circular. Its falsifier is correspondingly concrete: a memory-bearing system that cannot be decomposed into these parts would end it.

After diagnostics comes architecture. It is not enough to know the ledger, the selector, and the signatures; one must also know what persistent memory-bearing organization can actually implement them. Part VI therefore closes the Qi-4 block by deriving the architectural layer – MePMoS, NHT, HAL, BFFT, and the exploration-to-commitment switch – that makes the earlier pillars realizable in open systems with memory.

45.1 Historical background (with citations)

Antecedents: Hebb (1949), Hopfield (1982), Marr (1971), 't Hooft (1993), Susskind (1995), Maldacena (1998), O'Keefe (1971), Buzsáki (2006), Tononi (2004) (Friston, 2010, Hebb, 1949, Hopfield, 1982, Maldacena, 1998, McCulloch & Pitts, 1943, Susskind, 1995, 't Hooft, 1993, Tononi, 2004); Qi-HAL subsumes these as gauges or limiting cases of Theorem T-HAL.1 $\cup$ T-HAL.2 (the MePMoS/NHT/HAL/BFFT memory architecture is the boundary projection of the 6D bulk and its phase-holographic coding; see A1 Sec. 7 and A4 Sec. 6 for full history).

Dual-layer architecture and the compatibility relation $\mathfrak{C}_{\text{HAL}}$. Qi-HAL is, in the L16 language, a dual-layer principal pillar. Its *algebraic layer* (T-HAL.1) is the Qi-HAL universality theorem (Theorem 53.1): the named hierarchy MePMoS-NHT-HAL-BFFT is a sufficient boundary realization of the GRIP cascade on the family-projector + memory moduli, with the $U(1)_{Q_i}$ gauge sector as the memory carrier. Its *operational/physical layer* (T-HAL.2) is the body of empirical realizations: the Hopfield network (Hopfield, 1982), the place cells of O'Keefe & Dostrovsky (1971), the *Drosophila* ring attractor of Turner-Evans et al. (2020), and the thalamocortical forward-model architecture of Llinás 2001. The two layers are tied together by the named compatibility relation

$$\boxed{\mathfrak{C}_{\text{HAL}} \equiv \text{T-HAL.1} \implies \text{T-HAL.2} \quad (\text{via the } U(1)_{Q_i} \text{ memory carrier and the GRIP cascade on persistent memory moduli}).} \quad (45.1)$$

The universality theorem at the algebraic level forces any sufficiently persistent open system with informational currents to implement the five-layer hierarchy; the empirical realizations across neuroscience (place cells, ring attractors, thalamocortical loops) are then the boundary instances of this structural necessity.

46 Constitutional statement of the Qi-HAL principle

The architectural claim of Qi-HAL begins with a minimal observation: an open system counts as more than a passive dissipator only if part of its present state depends on past interactions in a way that remains functionally available for future action. This is the generic notion of memory, and it is intentionally broad enough to cover synapses, gene-regulatory marks, chemical modification states, conformational cycles, trail networks, field configurations, and externalized collective memory.

Definition 46.1 (Memory-bearing open system). An admissible open system with state packet

$$\Gamma(t) = (x^A(t), J^\alpha(t), \mathcal{P}, \mathcal{O})$$

is called *memory-bearing* if there exists an internal or external state variable $m(t) \in \mathcal{M}$ and a causal kernel $K_m(\tau)$ with support on $\tau \geq 0$ such that

$$m(t) = \int_0^\infty K_m(\tau) \Xi(x(t-\tau), s(t-\tau), a(t-\tau)) \, d\tau, \quad (46.1)$$




Figure 12: The architecture a system falls into once it keeps a trace for longer than a single cycle of acting and looking — the entrance condition at the top, which is a comparison of two times and nothing more. The loop’s order is not a modeling choice: sensing cannot come before action without the system knowing an observation it has not yet caused, and prediction cannot come after sensing without becoming hindsight, so each rearrangement destroys a step that has work to do. Three motifs hang off the loop. What is stored is a set of relations between phases rather than the phases themselves, which is why the same content survives being carried on a completely different rhythm and why a chemical, a neural and a digital system can hold it on unrelated hardware. The store is a lattice rather than a filing cabinet, which is why damage degrades a memory instead of deleting a list. And the ladder of nested timescales is only as deep as the spectral window of the preceding Part allows. At the foot: a write is an elementary boundary event, so the ledger of the first pillar prices it. What is established is a canonical normal form rather than a proof that no other shape could exist.

for some local update functional Ξ , and such that the associated retention time

$$\tau_m := \frac{\int_0^\infty \tau \|K_m(\tau)\| \, d\tau}{\int_0^\infty \|K_m(\tau)\| \, d\tau} \quad (46.2)$$

exceeds the characteristic observation-to-action cycle time T_c of the system:

$$\tau_m > T_c. \quad (46.3)$$

Definition 46.1 is intentionally operational. It does not assume that memory is neural, digital, or symbolic. It requires only that the system carry forward a compressed trace of earlier interactions long enough for that trace to matter. Equation (46.3) is the real entrance condition: without persistence beyond one cycle, memory cannot guide prediction or action.

The second ingredient is the informational gauge symmetry introduced earlier. In the present program a memory trace is not an abstract bookkeeping label added from outside; once $U(1)_{Qi}$ is present, informational events become physical current-carrying events. This promotes memory from a descriptive convenience to a gauge-bearing dynamical sector.

A note on why gauge structure is the right framework for memory. Any quantity that propagates in time and admits a phase ambiguity — and a memory trace is exactly such a quantity, since its absolute phase carries no observable content — must, if it is to be physical, be expressed in gauge-invariant relational form: as a phase *difference* corrected by holonomy. This is the standard reasoning that makes electromagnetic potentials gauge-invariant only modulo their gradient. Applied here, it makes phase-coded memory compatible with the Arisaka ledger of Qi-QUA. Readers from biology may think of this as the precise way to say that a memory’s content is what it differs from, not what it absolutely is; readers from neuroscience will recognize it as the formal version of the cocktail-party intuition that the relevant signal is always relative to a reference.

The constitutional data of the informational carrier are a real pair: a real density $\rho_Q \geq 0$, the local informational intensity, and a real, circle-valued phase ϕ_Q , an angle on the compact $U(1)_{Qi}$ memory circle. This pair is what the memory sector stores — a real angle on a real circle is the minimal memory register — and nothing complex is postulated to hold it. Because $U(1)_{Qi}$ is a compact bulk gauge circle (the four-component BI-6 sector F_{Qi} of Sec. 3.3), the circle’s quarter-turn is available as the unique circle-compatible complex structure on Qi -charged sectors — the derived gauge i of T-P4-1 of the Foundations Branch (Arisaka, 2026, P4-Real); the causal clock supplies the state-space i (T-P4-2), and the handshake forces the two to agree on shell (T-P4-3). Under that derived quarter-turn the real pair renders locally as the single complex carrier

$$\Psi_Q(x, t) = \sqrt{\rho_Q(x, t)} e^{i\phi_Q(x, t)}, \quad (46.4)$$

the Madelung packaging of (ρ_Q, ϕ_Q) : its i is the boundary’s shorthand for the derived quarter-turn, not a primitive, and its modulus and argument return the constitutional pair $(\sqrt{\rho_Q}, \phi_Q)$. The polar dressing is convenience; the reality beneath it is the real density and the real circle-valued phase. This is A1’s register discipline for the Hilbert space $\mathcal{H}$ applied to the memory sector: the complex scalars are boundary shorthand, priced at candidate grade by L-Constitution-1 and recorded in

the AX-3 register note of [Arisaka \(2026, A1-Constitution\)](#). A local $U(1)_{Q_i}$ gauge transformation with parameter $\Lambda(x, t)$ is then a real shift of the stored angle, $\phi_Q \mapsto \phi_Q + q_Q \Lambda$, which the dressing renders as

$$\Psi_Q \mapsto e^{iq_Q \Lambda} \Psi_Q, \quad A_\mu^{(Q_i)} \mapsto A_\mu^{(Q_i)} + \partial_\mu \Lambda, \quad (46.5)$$

so the covariant derivative, whose i is again the derived quarter-turn, is

$$D_\mu \Psi_Q = (\partial_\mu - iq_Q A_\mu^{(Q_i)}) \Psi_Q. \quad (46.6)$$

The crucial consequence is that the physically meaningful memory token cannot be an absolute phase. It must be a relational, gauge-invariant quantity. The minimal such token is a phase difference corrected by the gauge holonomy:

$$\Phi_{ij}^{\text{inv}} = \phi_i - \phi_j - q_Q \int_{\gamma_{ij}} A_\mu^{(Q_i)} dx^\mu, \quad (46.7)$$

where γ_{ij} joins two encoded states or channels. In particular, around a closed loop γ one obtains the informational holonomy

$$\mathcal{W}_\gamma^{(Q_i)} = \exp \left(iq_Q \oint_\gamma A_\mu^{(Q_i)} dx^\mu \right). \quad (46.8)$$

Equation (46.7) is the first real bridge from Qi-QUA and Qi-EP to Qi-HAL. Once memory is carried by a gauged informational field, relative phase and holonomy become natural state variables. That is why a memory-bearing open system under $U(1)_{Q_i}$ can be re-expressed in an oscillator- and phase-based architecture even when its microscopic substrate is chemical, mechanical, biological, or social.

Definition 46.2 (Qi-HAL). Qi-HAL is the canonical architectural factorization of a memory-bearing, $U(1)_{Q_i}$ -carrying open system into five mutually coupled layers:

- (i) **MePMoS**: the causal loop of Memory $\rightarrow$ Prediction $\rightarrow$ Motion $\rightarrow$ Sensing $\rightarrow$ updated Memory;
- (ii) **NHT**: a phase-holographic encoder that maps latent geometry or latent state to observable oscillatory phases and delays;
- (iii) **HAL**: a gauge-invariant memory lattice storing relational phase structure and attractor couplings;
- (iv) **BFFT**: a dyadic multiscale hierarchy coordinating slow/global and fast/local processing;
- (v) **OH+IH=EIC**: the internality layer — the outer hologram of externally registered boundary data, the inner hologram of the internally maintained model, and the closed loop $\text{EIC} := \mathcal{C}_{\text{loop}}(\text{OH}, \text{IH})$ that couples them.

Remark 46.3 (Why the internality layer belongs inside Qi-HAL, and the switch does not). Two membership decisions are worth stating, because the earlier formulation made the opposite pair. $\text{OH} + \text{IH} = \text{EIC}$ belongs *inside* the definition: a memory-bearing system does not merely store and recall, it maintains a model of the environment it is embedded in and acts through that model,

and every application in Part IX and Part X — chemistry, life, brains, collectives — turns on that loop rather than on the four storage-and-timing layers alone. Leaving it outside made Qi-HAL an architecture of bookkeeping and put the physics one part away.

The DGI–GSSB switch, by contrast, does not belong in an GG-A7-Qi4 definition at all. It is GG-A6-GRIP’s runtime, imported here and not restated: this paper consumes the cascade, it does not define it. Nor was it doing work as a layer — the present section’s own honesty note (Remark 53.2) records that its necessity argument was circular, restating the entrance condition rather than deriving anything from it. Removing it costs nothing and removes a restatement of upstream content from a downstream definition.

Definition 46.2 is intentionally architectural rather than phenomenological. It does not say that all systems look identical. Rather, it says that once persistent memory and informational gauge current are present, the runtime can be decomposed into these five tasks. Some systems implement them minimally or externally; some implement them richly and internally. The job of the rest of the section is to make that decomposition exact.

Remark 46.4 (Why the word “neural” in NHT is historical rather than restrictive). The acronym NHT is retained because the mature experimental examples come from neural systems. In the present paper, however, the mathematics is not restricted to neurons. Any substrate that converts latent geometry into phase, delay, or oscillatory code qualifies. Chemical oscillators, conformational cycles, trail networks, communication rhythms, and digital clocked systems can all instantiate the same structure.

Table 33: The five architectural motifs of Qi-HAL. Each row is one motif, with the question it answers, its primary mathematical object, and its operational role. Read downward the five compose into a loop: memory feeds prediction, prediction feeds motion, motion feeds sensing, sensing updates memory, and the last row is the switch that lets a system reorganize its own landscape rather than only descend on it. That last capability is what distinguishes a memory-bearing system from a merely open one.

Component	Question answered	Primary mathematical object	Operational role
MePMoS	How is one cycle of action and update closed?	Causal loop map $m_n \mapsto \hat{y}_n \mapsto a_n \mapsto s_{n+1} \mapsto m_{n+1}$	Organizes prediction, action, observation, and memory update.
NHT	How is latent geometry encoded?	Gauge-corrected phase map $y \mapsto \phi(y)$	Converts hidden multidimensional state into observable phases and delays.
HAL	Where is the memory stored?	Coupling matrix/tensor W and attractor energy E_W	Stores and stabilizes phase relations or structural traces.
BFFT	How are scales coordinated?	Dyadic ladder $f_n = 2^n f_0$ and cross-frequency couplings	Couples slow/global predictions to fast/local details.
DGI–GSSB	How does the system switch between exploration and commitment?	Symmetry-preserving orbit and broken-symmetry potential	Governs attention, search, selection, and decision.

47 Why persistent memory plus $U(1)_{Qi}$ is the true entrance condition

The user-level intuition is simple: once a system has memory, prediction becomes possible; once prediction becomes possible, action becomes meaningful; once action becomes meaningful, feedback must be compared to prediction; and once that comparison is repeated, a stable architecture is required. The mathematics of Qi-HAL simply packages this logic in a gauge-covariant way.

The argument becomes sharper when written in the language of the earlier sections. In Part III we introduced the current ledger and the Arisaka Equation,

$$\frac{J_E}{Q_B} = J_e = J_B = J_Q.$$

A memory write therefore cannot be free. If a process creates one additional persistent informational event, there is a corresponding minimal denominator in the energy channel. This does not yet say how the event is stored, but it already forbids the fiction that memory is merely symbolic. Memory consumes budget.

In Part IV we showed that selected trajectories, fluxes, and posterior states arise from a common projected extremal principle. A memory variable becomes useful only when it participates in that selection, namely when it changes which path is chosen, which current is sustained, or which posterior is stabilized. Memory without selection is inert archival residue.

In Part V we then showed that persistent open systems generically develop a scale-covariant window with broad fluctuations and two-knee spectra. A memory architecture operating in such a window must coordinate multiple timescales. This is exactly the role later assigned to BFFT.

Taken together, the previous sections imply the following entrance proposition.

Proposition 47.1 (Architectural entrance proposition). *Let an admissible open system satisfy the following four conditions:*

- (a) *it is memory-bearing in the sense of Definition 46.1;*
- (b) *it carries nonzero informational current J_Q coupled to $U(1)_{Qi}$;*
- (c) *its actuation changes future observations through at least one open port; and*
- (d) *its selected runtime contains a nontrivial Qi-CON window or, at minimum, more than one dynamically relevant timescale.*

Then the runtime admits a nontrivial architectural factorization of the form given in Definition 46.2.

Proof. Condition (a) supplies a history-bearing state variable m . Condition (b) permits the informational state to be written in gauge-covariant amplitude-and-phase form (46.4)–(46.7), thereby making relational phase storage natural. Condition (c) ensures that memory can affect future input through action; this closes the MePMoS loop. Condition (d) guarantees that at least a two-level scale protocol exists, allowing slow and fast channels to separate. Once these four ingredients are present, one can define the causal loop (MePMoS), the phase representation (NHT), the persistent

coupling substrate (HAL), the multiscale ladder (BFFT), and the exploration/commitment switch (DGI-GSSB). The construction of each layer is given explicitly below. ■

Proposition 47.1 is deliberately modest. It does not say that every such system has a large brain-like architecture. It says that the *normal form* of its runtime is architectural rather than purely differential-equation-like. Minimal systems realize the five motifs in degenerate or compressed form; richer systems realize them in distributed, nested, and highly nonlinear form.

48 MePMoS: derivation of the causal loop

We begin with the most elementary question: what is the unique causal ordering of memory, prediction, action, and sensing in a system that both acts and learns? In plain language the answer seems obvious, but in a theorem-driven paper it is worth stating carefully because the order is often blurred when perception and action are written as simultaneous updates.

Definition 48.1 (MePMoS). For a memory-bearing open system, MePMoS is the minimal causal loop

$$\mathcal{M}_n \xrightarrow{\mathcal{P}} \hat{y}_n \xrightarrow{\mathcal{A}} a_n \xrightarrow{\mathcal{S}} s_{n+1} \xrightarrow{\mathcal{U}} \mathcal{M}_{n+1}, \quad (48.1)$$

where $\mathcal{M}_n$ is the memory state available at the beginning of the cycle, $\hat{y}_n$ a predicted latent or sensory outcome, a_n the action or motion command, s_{n+1} the next sensed input, and $\mathcal{U}$ the memory update map.

The essential point is that the system cannot directly map old memory to new memory if it is to remain adaptive; it must pass through a prediction step and an action step. Prediction alone is not enough, because without action there is no way to test the prediction by changing the environment or the vantage point. Likewise sensing cannot precede motion in a closed action-perception loop because the relevant new sensation is partly determined by the very action chosen in the current cycle.

Theorem 48.2 (Causal uniqueness of MePMoS (T-Qi4-12.6)). *Consider a discrete open-system cycle with four macro-operations: memory readout, prediction, action, and sensing. Assume (i) prediction may depend only on information already stored at the start of the cycle, (ii) action may depend on the prediction and previously stored context, (iii) the incoming sensory record at the end of the cycle may depend on the chosen action and the environment, and (iv) the updated memory at the end of the cycle may depend on the old memory, the prediction, and the realized sensory record. Then, up to trivial relabeling inside each macro-operation, the only admissible topological order is*

$$\text{Memory} \rightarrow \text{Prediction} \rightarrow \text{Motion} \rightarrow \text{Sensing} \rightarrow \text{updated Memory}.$$

Proof. The assumptions define a dependency graph with arrows

$$\mathcal{M}_n \rightarrow \hat{y}_n, \quad (\mathcal{M}_n, \hat{y}_n) \rightarrow a_n, \quad a_n \rightarrow s_{n+1}, \quad (\mathcal{M}_n, \hat{y}_n, s_{n+1}) \rightarrow \mathcal{M}_{n+1}.$$

Any topological sort must therefore place $\mathcal{M}_n$ before $\hat{y}_n$, $\hat{y}_n$ before a_n , a_n before s_{n+1} , and s_{n+1} before $\mathcal{M}_{n+1}$. No permutation violating these arrows is admissible: placing sensing before action would require the system to know a future observation before choosing the action that partly generates it; placing prediction after sensing would collapse prediction into postdiction; placing memory update before new sensing would erase the error-correction step. Hence the macro-order is unique. ■

The discrete cycle can be written more concretely as

$$\hat{y}_n = \mathcal{P}(\mathcal{M}_n), \quad (48.2a)$$

$$a_n = \mathcal{A}(\hat{y}_n, \mathcal{M}_n), \quad (48.2b)$$

$$s_{n+1} = \mathcal{S}(a_n, \eta_{n+1}), \quad (48.2c)$$

$$\mathcal{M}_{n+1} = \mathcal{U}(\mathcal{M}_n, \hat{y}_n, s_{n+1}). \quad (48.2d)$$

Here η_{n+1} denotes environmental input. The operational error signal is

$$\epsilon_{n+1} = s_{n+1} - \hat{s}_{n+1|n}, \quad (48.3)$$

where $\hat{s}_{n+1|n}$ is the sensory consequence predicted at stage n . In an inferential implementation one may take

$$\mathcal{M}_{n+1} = \mathcal{M}_n - \eta_m G_m^{-1} \nabla_{\mathcal{M}} \left[\frac{1}{2} \|\epsilon_{n+1}\|^2 + \lambda_R R(\mathcal{M}_n) \right] + \xi_n, \quad (48.4)$$

with G_m a metric on memory space, R a regularizer or prior cost, η_m a learning step, and ξ_n a fluctuation term. Equation (48.4) is not a new principle added by hand. It is simply the Qi-EP/FEPTYPE selector written in the memory coordinates of the architecture layer (Friston, 2010).

In continuous time, writing the memory state as $m(t)$ and the predictive free-energy-like cost as $\mathcal{F}(m, s, a)$, one obtains the normal form

$$\tau_m \dot{m} = -\frac{\delta \mathcal{F}}{\delta m} + \chi_Q J_Q + \zeta(t), \quad (48.5)$$

where the explicit J_Q term reminds us that memory maintenance is an informational-current-bearing process rather than an abstract computational update.

The interpretation is extremely general. In a molecular machine the memory state may be a conformational coordinate or chemical occupancy. In a bacterium it may be a receptor methylation state. In a mammalian brain it may be a large distributed recurrent state. In a society it may be a distributed archive coupled to communication channels. The causal grammar is unchanged.

Table 34: The MePMoS loop across domains. Four systems run down the first column, from a molecular machine to a collective; the remaining columns identify what plays the part of memory, prediction, motion and sensing in each. The table’s claim is that the loop is the same loop in all four rows and that only the substrate changes — which is what makes the motif a structural statement rather than a biological one. Rows are ordered by increasing spatial scale, not by increasing sophistication.

Domain	Memory	Prediction	Motion/action	Sensing/update
Molecular machine	Conformational occupancy	Expected next catalytic state	Step, rotation, or gating bias	Product release, ligand binding, torque feedback
Cell or microbe	Internal biochemical state	Expected gradient or nutrient change	Chemotactic step or regulatory action	Receptor occupancy, signal update
Nervous system	Recurrent network state	Predicted sensory or latent state	Motor command or attentional shift	Sensory stream, error correction
Collective system	Shared archive, communication state	Expected collective outcome	Policy, message, or institutional action	External response, archive revision

49 NHT: phase-holographic coding of latent geometry

Once a causal loop exists, the next question is how latent structure is encoded so that it can be stored, compared, and updated efficiently. The answer proposed by Qi-HAL is phase-holographic coding. Again the word “neural” in NHT is historical; the derivation is not specifically neural.

Suppose the relevant latent variable of the system is $y = (y^1, \dots, y^d) \in \mathcal{Y}$. A minimal phase encoding is

$$\phi_i(t) = \omega_i t + \frac{2\pi}{L_i} y^i(t) + \lambda(t), \quad i = 1, \dots, d, \quad (49.1)$$

where ω_i is a carrier frequency, L_i a characteristic scale, and $\lambda(t)$ a common gauge phase. Under $U(1)_{Q_i}$ the physically meaningful coordinate is the gauge-corrected phase

$$\tilde{\phi}_i(t) = \phi_i(t) - q_Q \int^t A_{0,i}^{(Q_i)}(\tau) d\tau. \quad (49.2)$$

Therefore differences in latent state can be reconstructed from phase differences:

$$\delta y^i = \frac{L_i}{2\pi} \delta \tilde{\phi}_i. \quad (49.3)$$

This immediately gives the local induced metric

$$ds_{\mathcal{Y}}^2 = \sum_{i=1}^d \left(\frac{L_i}{2\pi} \right)^2 (d\tilde{\phi}_i)^2. \quad (49.4)$$

Proposition 49.1 (Metric recovery by phase tomography). *If the latent coordinates are encoded according to Eq. (49.1) and read out through the gauge-corrected phases (49.2), then the local Euclidean metric on the latent chart is recovered exactly by Eq. (49.4). More generally, a curved*

latent metric $g_{ij}(y)$ may be encoded by replacing the diagonal scale factors $(L_i/2\pi)^2$ with a positive-definite phase metric tensor $G_{ij}(\phi)$.

Proof. Differentiate Eq. (49.1) at fixed carrier-time slice or, equivalently, subtract the trivial carrier contribution $\omega_i dt$. Then $d\tilde{\phi}_i = (2\pi/L_i)dy^i$. Multiplying by $(L_i/2\pi)^2$ and summing over i yields Eq. (49.4). The curved case follows by replacing the diagonal quadratic form with a general positive-definite one. ■

The conceptual content of Eq. (49.4) is far-reaching. A high-dimensional latent world can be encoded in relative phases, delays, or rhythmic offsets. The system does not need a literal geometric map drawn somewhere inside it. It only needs a phase configuration from which distances, alignments, or transition costs can be reconstructed.

A familiar special case is a ring attractor, in which one coordinate $y \in S^1$ is encoded as a phase angle ϕ . But the same logic extends to toroidal, grid-like, chemical, and distributed memory systems. In a reaction network, phase may mean oscillation phase of concentrations. In a mechanical rotor, it may mean angular coordinate. In a social rhythm, it may mean synchronization offset relative to a shared clock.

The holographic language is justified because the code is not a point-by-point copy of the latent world. Rather, the latent geometry is stored distributively in phase relations, just as ordinary holography stores spatial information distributively in interference structure. This is why a damaged or partial substrate may still retain partial reconstructive power.

Table 35: Interpretations of NHT beyond neuroscience. Four substrates run down the first column, with the variable each encodes, the phase-like carrier that does the encoding, and what a tomographic readout would mean in each case. The table is here because the motif’s name comes from neuroscience and its content does not: a chemical oscillator network encodes state in relative phase for the same structural reason a neural population does, and confining the motif to brains would be a fact about the name rather than about the mathematics.

Substrate	Encoded variable	Phase-like carrier	Meaning of tomography
Chemical oscillator network	Concentration or environmental phase	Oscillation phase of coupled reactions	Recover latent chemical/environmental coordinates from relative phase
Conformational molecular cycle	Position on catalytic cycle	Rotor angle or cycle phase	Reconstruct step location and expected next state
Neural population	Position, heading, or latent cause	Spike phase / local field potential phase	Recover latent spatial or causal geometry
Collective communication system	Social or task state	Synchronization phase across agents/clocks	Recover distributed shared state from timing relations

50 HAL: the holographic attractor lattice as a memory substrate

Phase coding alone is not memory. A system also needs a substrate in which phase relations can leave durable traces. HAL is that substrate. In a minimal realization it is simply a set of couplings that stores relational structure. In a rich realization it becomes a high-dimensional attractor lattice supporting recall, completion, interpolation, and stable world-model reconstruction.

Let the encoded phase-bearing features be

$$\sigma_i^{(\alpha)}(n) = \exp(i\phi_i^{(\alpha)}(n)), \quad (50.1)$$

where i labels local units or channels and α labels feature families or subbands. A minimal write rule is the gauge-invariant correlator update

$$W_{ij}^{\alpha\beta}(n+1) = (1 - \lambda_W)W_{ij}^{\alpha\beta}(n) + \gamma_W \sigma_i^{(\alpha)}(n) \overline{\sigma_j^{(\beta)}(n)}, \quad (50.2)$$

with decay or forgetting rate λ_W and write strength γ_W . Under a common local $U(1)_{Q_i}$ phase shift, the product $\sigma_i^{(\alpha)} \overline{\sigma_j^{(\beta)}}$ is invariant whenever the two channels carry the same informational charge. In the more general case one inserts the parallel-transport factor of Eq. (46.7); the same conclusion follows.

The corresponding readout energy is

$$E_W[\sigma] = -\frac{1}{2} \sum_{\alpha,\beta} \sum_{i,j} W_{ij}^{\alpha\beta} \sigma_i^{(\alpha)} \overline{\sigma_j^{(\beta)}} + \sum_{\alpha,i} U(|\sigma_i^{(\alpha)}|^2), \quad (50.3)$$

with a local confining term U . Stable memories are local minima of E_W . When the couplings are organized on a ring, torus, graph, membrane, or more general lattice, those minima define a geometric memory manifold. That manifold is what the name holographic attractor lattice denotes.

The connection to the earlier pillars is immediate. A write event changes W , and therefore costs energy according to the Qi-QUA ledger. A read event queries one or more minima of E_W , and therefore participates in Qi-EP trajectory or posterior selection. The existence of a whole family of minima across scales is precisely what Qi-CON diagnosed indirectly through broad spectral and spatial statistics.

The historical analogy to associative memory is useful but incomplete (Hopfield, 1982). In ordinary associative memory one stores patterns. In HAL one stores gauge-covariant relations, phase geometry, and attractor structure. This makes HAL more naturally compatible with continuous variables, manifold-valued latent states, and topological phases than a purely binary memory model.

Proposition 50.1 (Gauge-invariant memory storage). *If the write rule is constructed from phase-difference correlators of the form in Eq. (50.2), optionally with explicit parallel transport for channels living in different gauge patches, then the stored memory tensor W is invariant under local $U(1)_{Q_i}$ gauge transformations. Consequently the physically stored object is relational phase geometry rather than absolute phase.*

Proof. Under $\sigma_i^{(\alpha)} \mapsto e^{iq_Q \Lambda_i} \sigma_i^{(\alpha)}$, the factor $\sigma_i^{(\alpha)} \overline{\sigma_j^{(\beta)}}$ acquires the phase $e^{iq_Q (\Lambda_i - \Lambda_j)}$. If the channels are in the same local patch and the same charge sector, this phase cancels for equal local gauges or is absorbed into the covariant phase-difference variable. More generally, inserting the Wilson-line transport factor between i and j cancels the residual phase exactly. Therefore the stored object depends only on relational gauge-invariant data. ■

The most important practical question is whether HAL depends on absolute frequency. If a memory is written at one carrier frequency and read at another, is the content lost? The answer in the idealized relational limit is no.

Theorem 50.2 (Frequency-independent HAL readout (T-Qi4-12.9)). *Suppose a pattern is written using oscillatory carriers*

$$x_i^w(t) = \exp(i(\omega_w t + \phi_i)), \quad (50.4)$$

and stored via the ideal correlator

$$W_{ij} \propto \left\langle x_i^w(t) \overline{x_j^w(t)} \right\rangle_t = e^{i(\phi_i - \phi_j)}. \quad (50.5)$$

If the pattern is later queried using carriers

$$x_i^r(t) = \exp(i(\omega_r t + \psi_i)), \quad (50.6)$$

then the overlap energy

$$\mathcal{O}[\psi|W] = \sum_{i,j} \overline{x_i^r(t)} W_{ij} x_j^r(t) \quad (50.7)$$

is independent of the common read frequency ω_r . It depends only on the relative phase pattern ψ_i compared to the stored relative phase pattern ϕ_i .

Proof. Substituting Eqs. (50.5) and (50.6) into Eq. (50.7) yields

$$\mathcal{O}[\psi|W] \propto \sum_{i,j} e^{-i(\omega_r t + \psi_i)} e^{i(\phi_i - \phi_j)} e^{i(\omega_r t + \psi_j)} = \sum_{i,j} e^{i[(\phi_i - \phi_j) - (\psi_i - \psi_j)]}.$$

The carrier phases $\omega_r t$ cancel identically. Hence the overlap depends only on relative phase differences. The common read frequency therefore does not affect ideal recall. ■

Theorem 50.2 has two important consequences. First, HAL is fundamentally a relational storage principle. Second, it can be realized across very different substrates. A chemical system, a neural system, and a digital clocked system may use entirely different carrier frequencies, yet the same relational content can in principle be preserved.

One can also define a simple architectural completeness index for this pillar alone — written QCI_{HAL} , to keep it distinct from the four-pillar Qi-4 Completeness Index of Part VII —

$$\text{QCI}_{\text{HAL}} = \frac{1}{5} (I_{\text{MePMoS}} + I_{\text{NHT}} + I_{\text{HAL}} + I_{\text{BFFT}} + I_{\text{switch}}), \quad (50.8)$$

where each indicator $I \in [0, 1]$ measures the degree to which the corresponding Qi-HAL component is realized in a given substrate. Minimal biochemical toggles may have $\text{QCI}_{\text{HAL}} \ll 1$, while distributed nervous systems may have $\text{QCI}_{\text{HAL}} \rightarrow 1$. Equation (50.8) is not needed for the theory itself, but it is useful as a future empirical scoring tool.

51 BFFT: the dyadic multiscale protocol

The moment memory becomes nontrivial, one also needs scale coordination. Not every variable changes at the same rate; not every prediction is local; not every correction is fast. The natural multiscale protocol in Qi-HAL is BFFT, the binary fractal Fourier transform. The name is meant broadly: one may realize it dynamically by a nested oscillator hierarchy or analytically by a dyadic wavelet/Fourier-packet decomposition (Mallat, 2009). The key point is the dyadic scale ladder.

The derivation begins from the conformal radial coordinate already used in GG-6 and later inherited by Qi-CON,

$$r = \ln \left(\frac{\mu_0}{\mu} \right). \quad (51.1)$$

Equal increments in r correspond to multiplicative steps in physical scale. Suppose a scale hierarchy is sampled by a constant increment $\Delta r = \ln \lambda$, so that

$$\mu_{n+1} = \lambda \mu_n, \quad f_{n+1} = \lambda f_n. \quad (51.2)$$

If one also demands integer nesting of cycles between adjacent levels, the multiplicative factor λ must be an integer greater than one. The minimal nontrivial choice is $\lambda = 2$, and this is exactly the binary or dyadic ladder. Measured neural cross-frequency coupling is, however, typically *non*-dyadic — the hippocampal theta–gamma code nests at roughly 1:5–1:8 (Lisman & Jensen, 2013), an integer $m \approx 7$; the dyadic $\lambda = 2$ is therefore the minimal idealization, and the empirical ladder selects the architecture-specific integer $m > 1$, not a prediction that biological coupling is binary.

Proposition 51.1 (Dyadic minimality). *Among multiplicative scale hierarchies with the properties (i) exact parent-child nesting of periods and (ii) the minimal number of additional structural choices beyond the constitutional ledger, the unique minimal nontrivial ratio is*

$$\lambda = 2. \quad (51.3)$$

Hence the natural Qi-HAL scale ladder is

$$f_n = 2^n f_0. \quad (51.4)$$

Proof. Exact parent-child nesting requires that one period at level n contain an integer number $m > 1$ of periods at level $n + 1$, hence $f_{n+1} = m f_n$ with integer m . The requirement of minimal extra structure chooses the smallest admissible integer, namely $m = 2$. Therefore $\lambda = 2$ and the dyadic ladder follows. ■

The dyadic ladder is not merely aesthetically simple. It is exactly the structure required to mediate top-down and bottom-up flow. Slow channels carry coarse, global context; fast channels carry fine, local details. A phase-locking action for adjacent levels is

$$S_{\text{lock}} = \sum_n \kappa_n \int dt [1 - \cos(2\phi_n - \phi_{n+1})]. \quad (51.5)$$

The stationary condition is

$$2\phi_n - \phi_{n+1} = \text{const}, \quad (51.6)$$

which is the canonical 1:2 phase-lock relation between neighboring scales.

This immediately gives a bridge to the two-knee Qi-CON spectrum. If the active conformal window spans from f_{IR} to f_{UV} , then the number of effectively available dyadic levels is approximately

$$N_{\text{lev}} \approx 1 + \left\lfloor \log_2 \left(\frac{f_{\text{UV}}}{f_{\text{IR}}} \right) \right\rfloor. \quad (51.7)$$

Thus the diagnostic spectral window fixes the architectural depth of the BFFT ladder. This is a concrete, falsifiable bridge between Pillar III and Pillar IV.

Scope note on the cosmological reading of BFFT. The companion GG-A10-Cosmos (Arisaka, 2026, A10-Cosmos) carries a proposed identification of the CMB damping multipole with a BFFT ultraviolet knee, $\ell_D = K(2\pi)^2 \approx 1382$; that reading is an *interpretive overlay*, sitting +2.3% above the framework’s own Boltzmann-computed damping scale ($\ell_D = 1350 \pm 3$), and is scored accordingly there. The present paper supplies no holonomy-potential normalization that would close that gap — the BFFT machinery here is information-architectural — so any promotion of the overlay must come from a bus normalization not yet constructed, not from this paper by citation.

For signal-analysis purposes one may write the same ladder as a dyadic transform,

$$X_n(\tau) = 2^{-n/2} \int x(t) \psi\left(\frac{t - \tau}{2^n}\right) dt, \quad (51.8)$$

which simply restates that the architecture can be understood either dynamically or analytically. In both readings the binary hierarchy is doing the same job: separating scales while preserving cross-scale couplings.

Table 36: Functional interpretation of the BFFT hierarchy. Three level types — slow and coarse, intermediate, fast and fine — run down the first column, with the characteristic scale, the architectural role and the relation to neighboring levels alongside. The last column carries the content: the levels are not independent channels but a dyadic ladder, so a change at one scale propagates in a definite direction, and the table says which.

Level type	Characteristic scale	Typical architectural role	Relation to neighboring levels
Slow / coarse	Low frequency, long horizon	Global context, priors, slow homeostatic set-points	Constrains the envelope or admissible region for faster levels
Intermediate	Mesoscale bands	Working coordination, routing, synchronization	Mediates bidirectional exchange between slow and fast levels
Fast / fine	High frequency, short horizon	Local detail, error correction, rapid sensory update	Reports deviations upward and refines lower-resolution predictions

52 The GRIP-loop as the exploration-to-commitment switch

The fifth architectural motif is not an additional physical object but a runtime mode switch. A memory-bearing system must alternate between two regimes. In one regime it explores, samples, and remains flexible across a set of near-equivalent possibilities. In the other regime it commits to one branch, one action, one interpretation, or one attractor. The first is naturally described by DGI; the second by GSSB. The stabilized result is GA.

A minimal normal form uses an order parameter m (which may be scalar or vector) and an effective potential

$$V_{\text{eff}}(m; t) = \frac{1}{2}a(t) \|m\|^2 + \frac{b}{4} \|m\|^4 - h(t) \cdot m, \quad b > 0. \quad (52.1)$$

When $a(t) > 0$, the symmetric point $m = 0$ is the unique minimum; the system remains exploratory. When $a(t) < 0$, a symmetry-broken shell or set of minima appears. The field $h(t)$ or small fluctuations then select one branch. That branch is stabilized by the quartic term and by environmental feedback, producing GA.

To connect this to observation and prediction error, define an evidence integral

$$I(t) = \int_{-\infty}^t \exp\left(-\frac{t-\tau}{\tau_I}\right) (\epsilon_\star^2 - \|\epsilon(\tau)\|^2) d\tau, \quad (52.2)$$

where $\epsilon(t)$ is the current prediction error and $\epsilon_\star$ is a reference tolerance. Taking

$$a(t) = a_0 - \lambda_I I(t) \quad (52.3)$$

shows that sufficiently accumulated evidence can drive $a(t)$ through zero. The system then moves from exploratory DGI to symmetry-broken commitment. If the environment later changes so that

error rises again, $I(t)$ decreases, $a(t)$ can return positive, and the system re-enters exploration. This is the normal form of attention, search, decision, commitment, and revision across many domains.

Proposition 52.1 (Exploration-to-commitment bifurcation). *For the potential (52.1) with $a(t)$ given by Eq. (52.3), there exists a critical evidence level*

$$I_c = \frac{a_0}{\lambda_I} \quad (52.4)$$

such that the system is in the exploratory DGI regime for $I < I_c$ and in the broken-symmetry commitment regime for $I > I_c$, up to perturbations induced by $h(t)$ and noise.

Proof. The sign change in $a(t)$ occurs exactly at $I = I_c$. For $a > 0$, V_{eff} has a single minimum at $m = 0$. For $a < 0$, the quartic term creates nonzero minima with radius $\sqrt{-a/b}$ when $h = 0$, or nearby perturbed minima for small $h \neq 0$. The system therefore crosses from a single exploratory basin to multiple committed branches at the critical evidence level. ■

This proposition turns a qualitative story into a concrete architectural rule. It says that a memory-bearing open system need not choose between flexibility and stability once and for all. Rather, it moves between the two in a controlled, evidence-dependent way. That is exactly what one expects of adaptive organisms, decision circuits, and collective institutions.

The GRIP-loop as the internalized regime switch. The two-regime vocabulary of GG-A6-GRIP (§7) makes precise what kind of object this switch is. When $a(t)$ crosses zero, the control parameters of the effective landscape (52.1) change and the basin structure of V_{eff} is reorganized — by definition a GRIP-Out event; the subsequent settling into the committed branch is GRIP-In descent; and the evidence dynamics (52.3) can run the cycle again. The GRIP-loop is therefore an *internally generated, recoverable alternation* (GRIP-Out · GRIP-In)*, and the contrast with the rest of the series is the discovery: the cosmological GRIP-Out’s of Arisaka (2026, A10-Cosmos) are exogenous and unrepeatable (the universe does not re-run its inflationary exit), whereas a memory-bearing SeOS carries its own landscape-reorganizing dial and schedules its reorganizations on evidence. This is precisely the microscopic mechanism that the capstone’s Cascade-Self-Sustenance corollary (GG-A11-Origin (Arisaka, 2026, A11-Origin)) requires: GRIP-Out events at the higher cascade layers are manufactured by the layers below — and Proposition 52.1 exhibits the minimal machine that manufactures them. One boundary clarification carries over from GG-A6-GRIP: ordinary port-driven excitation within a fixed landscape is *not* a GRIP-Out; the switch qualifies because $a(t)$ changes the landscape itself, not merely the system’s position on it.

53 A unified effective action for Qi-HAL

The preceding subsections can be compressed into a single effective architectural action. The point is not that every real system literally writes its microphysics in this form, but that every admissible implementation has this as a normal-form description after coarse-graining.

Let the dynamic variables be memory m , phase coordinates ϕ_i , memory couplings W , actions a , and sensed inputs s . Then an effective action is

$$\begin{aligned}
S_{\text{Qi-HAL}} = \int dt & \left[\underbrace{\frac{1}{2} \dot{m}^\top G_m \dot{m} + \frac{\lambda_\epsilon}{2} \|\epsilon\|^2}_{S_{\text{MePMoS}}} + \underbrace{\frac{1}{2} \sum_i \kappa_i \left(D_t \phi_i - \frac{2\pi}{L_i} \dot{y}^i \right)^2}_{S_{\text{NHT}}} \right. \\
& + \underbrace{\frac{\tau_W}{2} \|\dot{W}\|^2 + V_H(W, \sigma)}_{S_{\text{HAL}}} + \underbrace{\sum_n \kappa_n [1 - \cos(2\phi_n - \phi_{n+1})]}_{S_{\text{BFFT}}} \\
& \left. + \underbrace{V_{\text{eff}}(m; t)}_{S_{\text{switch}}} + q_Q A_\mu^{(Q_i)} J_Q^\mu \right]. \tag{53.1}
\end{aligned}$$

Here $D_t \phi_i = \dot{\phi}_i - q_Q A_{0,i}^{(Q_i)}$. The last term is the explicit informational gauge coupling. A natural partition or effective generating functional is then

$$\mathcal{Z}_{\text{Qi-HAL}} = \int \mathcal{D}m \mathcal{D}\phi \mathcal{D}W \mathcal{D}a \mathcal{D}s \exp \left[-\frac{S_{\text{Qi-HAL}}}{Q_B} \right]. \tag{53.2}$$

The appearance of the Arisaka Constant Q_B in the denominator is deliberate. It marks the ideal energetic denomination of an informational event, thereby tying the full architectural layer back to the current ledger of Qi-QUA.

With the first three Qi pillars already in place, the next theorem closes the Qi block by showing that the architectural layer is not a domain-specific add-on but the universal implementation class for memory-bearing open systems.

Theorem 53.1 (Qi-HAL canonical normal-form decomposition (T-Qi4-12.12)). *Let an admissible open system satisfy the hypotheses of Proposition 47.1. Then, after suitable coarse-graining to the timescales on which memory persists, its runtime admits an effective action of the form (53.1) and therefore a canonical factorization into MePMoS, NHT, HAL, BFFT, and the internality layer OH + IH = EIC. The decomposition is canonical in the sense that each motif corresponds to a structural requirement of the runtime, and the absence of any motif corresponds to a specific operational deficiency identified below; this is weaker than asserting that each motif is logically necessary in the strict sense (cf. Remark 53.2). The GRIP-cascade closure on which the factorization rests is, like its source in GG-A6-GRIP, conditional on that paper's remaining cascade conditioning (M11 in spectral-flow form; M10 resolved externally, the scope refinements closed at class level).*

Proof. We organize the five motifs into two groups: those that are structurally definitional (and therefore present by definition in any system that satisfies the hypotheses), and those whose presence is a non-trivial derived consequence. Sufficiency then follows from the canonical sum of sector actions.

Definitional motifs (MePMoS, OH + IH = EIC). By Definition 46.1, an admissible memory-bearing system possesses a memory coordinate whose update is a function of past sensing, prediction, and action; this closed loop is the structural definition of the MePMoS sector. The internality layer is definitional in the same sense and for the same reason: a system whose memory update

is action-conditioned is, by that very fact, maintaining an internal correlate of the environment it acts on and registering that environment at its boundary — which is what IH and OH name, and the action-conditioning is the closure of the loop between them. These two motifs are therefore present in any system that satisfies the hypotheses of the theorem; their presence carries no surplus content beyond the hypotheses themselves. Calling them “necessary” would be circular — they are necessary in the same sense that a triangle is necessarily three-sided. (What is *not* definitional, and is the substantive content of Part IX, is the structure of the loop once closed: the fidelity conditions, the maintenance ports, and the conditions under which the inner hologram tracks the outer one.)

Derived motif: NHT. This motif carries non-trivial content beyond the hypotheses. Under $U(1)_{Qi}$ gauge invariance, the physical observables that involve the memory coordinate must be gauge-invariant functions of m and its phase ϕ . The minimal gauge-invariant phase-coded observable on a non-simply-connected memory manifold is the Wilson holonomy of Eq. (46.8). Suppose a candidate system carries memory but realizes $U(1)_{Qi}$ trivially (so the holonomy is a pure-gauge constant). Then all phase content is unobservable, and the informational current J_{Qi} decouples from any measurable observable on the memory layer. This is a derived consequence, not a definition: it requires both the memory-bearing hypothesis *and* active $U(1)_{Qi}$ (cf. the L12 framing of $U(1)_{Qi}$ as constitutionally optional but empirically preferred), and it is operationally falsifiable by an experiment that demonstrates persistent memory in a system whose $U(1)_{Qi}$ sector is decoupled.

Derived motif: HAL. Persistence of memory traces over timescales long compared to the relaxation time of the boundary dynamics requires the existence of metastable basins in the effective potential of the memory coordinate. By the GA theorem (Theorem 7.2), such basins are exactly the attractors of the post-GSSB runtime; the collection of such attractors, organized by their gauge-invariant pairwise relations, is the HAL sector. Suppose the system has no metastable basins on the memory manifold: then traces decay on the boundary relaxation timescale, contradicting the operational definition of persistence (memory-bearing condition $\tau_m > T_c$, Eq. (46.3)). This is again a derived consequence (involving the GA theorem from Section 7) rather than a tautology.

Derived motif: BFFT. By Theorem 42.1, any memory-bearing system in the Qi-CON admissibility window exhibits a two-knee $1/f$ spectrum with a well-defined window between low-frequency knee f_L and high-frequency knee f_H . Reproducing such a spectrum requires scale-recursive structure with a geometric ratio between successive layers and binary commitment at each layer (a GSSB event resolving each indecision). Suppose the system has no nested binary scale structure: then either the $1/f$ spectrum cannot be reproduced (contradicting Theorem 42.1), or the binary-commitment structure of GSSB cannot be implemented (contradicting Theorem 7.3). This is derived from Qi-CON and the master pipeline rather than from the memory-bearing definition itself.

Sufficiency. Given the definitional motifs (MePMoS, OH + IH = EIC) and the three derived motifs (NHT, HAL, BFFT), the effective action of Eq. (53.1) is the canonical sum of their respective contributions: the memory potential S_m , the phase action S_ϕ , the Wilson-line action S_W , the action coordinate S_a , and the sensing coordinate S_s . Coarse-graining to the timescales on which memory persists bundles the sectors into the single effective action, and the decomposition follows from the canonical roles, not from the microscopic details of any particular substrate.

Thus the decomposition is sufficient, and three of the five motifs are derived consequences of the hypotheses; the remaining two are definitional. ■

Remark 53.2 (Honest scope of the canonical decomposition). It is tempting to advertise all five motifs as “logically necessary by contradiction,” but arguments of that shape are circular here: they restate the definition of the hypothesis rather than deriving the motif from it. The statement above therefore separates the two definitional motifs (MePMoS and the internality layer, both of which follow from the memory-bearing hypothesis itself) from the three genuinely derived ones, and states the result as a canonical normal-form decomposition rather than a strict universality claim. The empirical falsifier survives the change: a memory-bearing $U(1)_{Q_i}$ system that cannot be decomposed into the five sectors, or in which any one of the three derived motifs (NHT, HAL, BFFT) demonstrably fails to obtain, refutes Qi-HAL.

The theorem is the central result of this section. It makes precise a claim that is easy to state only loosely: once $U(1)_{Q_i}$ is applied, any open system with memory acquires a common architectural grammar. The grammar may be instantiated minimally in chemistry or richly in brains, but the runtime tasks are the same.

54 How Qi-HAL applies to any open system with memory

The universality claim becomes most persuasive when one resists the temptation to imagine only brains. The point of Qi-HAL is broader. Memory can be internal or external, analog or digital, molecular or social. The theorem does not require a cortex, a language, or even neurons. It requires only persistence, action-conditioned feedback, and informational gauge current.

Minimal biochemical case. A biochemical toggle, receptor cluster, or catalytic cycle may have a very small state space, yet if its present reaction bias depends on a stored modification state, it already satisfies Definition 46.1. The cycle phase plays the role of NHT the modified couplings or energy landscape play the role of HAL and slow/fast relaxation channels provide a minimal BFFT hierarchy. The DGI–GSSB switch appears when the system moves from a flexible undecided regime to a committed branch of reaction or conformation.

Cellular and microbial case. A chemotactic microbe or slime-mold network is an especially transparent open-system example (Nakagaki et al., 2000). Memory may be carried by receptor methylation, internal chemical state, or even an externalized trail network. Prediction appears as directional bias, motion as crawling or swimming, sensing as receptor update, and multiscale coordination as the nesting of fast local signaling within slow environmental adaptation. Here HAL may be partly externalized into the environment itself, showing that the theorem does not require memory to be stored in a single physically bounded organ.

Neural and cognitive case. In a nervous system the five motifs become more familiar. Recurrent state and synaptic structure supply MePMoS and HAL phase coding and population rhythms supply

NHT nested oscillatory bands or other multiscale population modes supply **BFFT** and evidence-driven selection produces the **DGI**–**GSSB** attention/decision cycle. The architecture is richer but not different in kind.

Collective and social case. A collective can also instantiate Qi-HAL. Shared archives, institutions, communication rhythms, distributed representations, and policy cycles provide memory, prediction, action, sensing, and update. The memory lattice is partly externalized in documents, infrastructures, or digital systems; yet the same logic applies. This is one of the reasons the present paper insists that the framework is for open systems in *all* sciences, not only for biology or physics.

Table 37: How the five Qi-HAL motifs appear across representative open systems with memory. Each row is one system, from a catalytic cycle to a collective; the columns take the motifs in turn and say what plays each part. Read across a row it is a description of one system; read down a column it is the claim that the motif is substrate-independent. The table repeats its header on every page because it is the reference a reader returns to when a later section names a motif without re-explaining it.

System	Memory	MePMoS	NHT	HAL	BFFT/switch
Catalytic cycle	Conformational occupancy, chemical mark	Cycle memory → catalytic bias → step → feedback	Rotor or oscillation phase	Modified energy landscape	Slow-fast nesting; branch commitment
Chemotactic cell	Internal adaptation, receptor memory, trail residue	Internal state → direction bias → movement → receptor update	Chemical oscillation phase	Internal network, possibly external trail	Adaptation bands; explore vs commit
Neural system	Recurrent state, synaptic matrix	Working memory → prediction → action → sensory update	Population phase code	Attractor manifold, synaptic lattice	Nested bands; decision bifurcation
Collective	Shared archive, institutions, digital storage	Archive → forecast → policy → measured outcome	Shared clocks, communication timing	Externalized memory lattice	Slow governance, fast signals; consensus

It is worth emphasizing the role of externalized memory. The theorem never assumed that memory must be stored behind the physical boundary of one organism. A beaver dam, a pheromone trail, a paper notebook, a digital server, or a laboratory instrument can all function as memory-bearing extensions of an open system. Once those extensions are coupled into the causal loop, they become part of the effective **HAL** sector. This is why the theory naturally extends from particles to organisms to societies.

Worked examples. ATP synthase as a minimal proto-Qi-HAL system, slime-mold/chemotactic collectives as externalized HAL, and ring-attractor cognitive maps as mammalian HAL are recorded in Appendix H.

55 Memory as something the world does

The habit of thinking about memory as storage is very hard to shake, and it is worth saying plainly what this Part replaces it with.

Storage suggests a place: a location where a thing is put and from which it is later retrieved, unchanged. Almost nothing about biological memory works that way, and the account here explains why it could not. What is held is not a set of items but a set of relations — differences between phases, not the phases themselves. That is why the same content survives being carried on a different rhythm entirely, so that a chemical system, a nervous system and a machine can hold the same thing on hardware with nothing in common. It is why damage degrades a memory rather than deleting a list. And it is why the question "where is it stored" has no good answer, not because the answer is hidden but because the question presupposes the wrong kind of object.

The second thing to carry away is that the loop's order is forced. A system cannot sense before it acts without knowing an observation it has not yet caused; it cannot predict after sensing without turning prediction into hindsight. There is only one way to arrange remembering, predicting, moving and sensing so that each step has something to do. Anything that keeps a past and uses it has to run the cycle in that order, whatever it is made of. That is a strong statement about architecture arrived at without looking at any particular architecture.

The third is the one this paper alone can make. The engine imported from the geometric companion needs something to reorganize its landscape from time to time, or a system would descend once and be finished. In cosmology those reorganizations arrive from outside and do not repeat. A system with a memory makes its own. It explores while its evidence is thin, commits when the evidence crosses a threshold, and goes back to exploring when the world changes and its predictions start to fail — and that cycle of attention, decision, commitment and revision is the same normal form in a bacterium choosing a direction, an animal choosing a route, and an institution choosing a policy. It is how something that would otherwise settle once keeps going.

With this Part the four pillars are complete, and it is worth standing far enough back to see what they are. The first says what a system spends and in what single currency. The second says which of the available ways of spending it the system takes. The third says what the taken path looks like to an instrument that changes the scale of its attention. The fourth says what the system has to be built like in order to do any of it a second time. Budget, selection, signature, architecture. That is not a list of four topics that happened to interest the same author. It is close to the minimal set of questions that must be answered before one is entitled to say that a thing has a way of going on rather than merely a way of happening. Answer the first three and you have described an event. Answer the fourth and you have described something with a history.

There is a temptation, having got this far, to say that these parts are logically necessary for anything that remembers. The Part declines to say it, and explains that the argument would be

circular — it would be restating the definition rather than deriving anything from it. What is offered instead is a normal form and a way to break it: find a memory-bearing system that cannot be decomposed this way, and the Part is finished. That is a smaller claim than it could have made, and a more useful one.

Two things the pillars name are left standing, and they are the subjects of what follows. The first is whether a system built this way in fact lasts: the pillars describe machinery, and machinery can run beautifully right up to the moment it stops, so persistence is a separate question and gets a separate answer in Part VII. The second is what happens when the loop just described is closed on itself — when a system’s picture of the world comes to include the system that is holding the picture. That is Part IX, and it is where an account of machinery turns into an account of a point of view.

Part VII

SeOS and QCI

The persistence class and the Qi-4 Completeness Index

Everything in nature is open. Almost nothing holds itself together. This Part is about the difference.

A speck of dust in a fluid is open — it exchanges momentum and energy with everything around it — but it maintains nothing. A convection cell is closer: it exists only while something keeps driving it, and it would be fair to say it is doing a kind of work to persist. A cell, an organism, a brain, an institution is a further step again, because what it maintains is not only a flow but an organization: a boundary, a set of records, a way of repairing itself when something goes wrong.

The first half of this Part turns that intuition into an inequality anybody can evaluate. What comes in through the ports, minus what is irreversibly lost, minus what repair costs, must stay positive on average over a window longer than the system's own repair time. Above zero, the system stays inside the region where it is still itself. Below zero for long enough, it leaves that region and stops being a member of the class. The margin is a computable quantity, not a judgment, and that is the whole point of writing it this way.

The second half asks a different question that is easy to confuse with the first. Surviving and running the full machinery of the previous four Parts are not the same thing. A chemical oscillator can persist indefinitely while realizing almost no memory architecture; a damaged nervous system can be rich in structure while sliding toward the edge of viability. So the Part introduces a second measure — how much of the four-pillar runtime a system actually runs — and deliberately builds it so that a missing pillar cannot be compensated by excellence in another. A system that scores nothing on one of the four scores nothing overall, because that is what incompleteness means.

Two things to hold on to while reading. The word used for the persistence class is a technical one borrowed from an older debate, and it carries no suggestion of intention or of moral standing; it means only that the local accounting is done against the survival of the packet under study. And the weights in the completeness measure are not derived. The paper says so. The honest use of that index is to compare two systems, not to certify one.

Part IX begins only after the full Qi stack is available. The next question is no longer how an open system is driven, selected, diagnosed, or architected, but which such systems persist and how fully they instantiate the Qi-4 runtime. Part VII therefore introduces the persistence class SeOS and the completeness diagnostic QCI.




Figure 13: Two questions this Part separates, which are usually asked as one. (a) Whether a thing can hold itself together is arithmetic: what comes in through the ports, minus what is irreversibly lost, minus what repair costs. The margin is a computable quantity rather than a judgment, and it can be evaluated for a protocell, a reef, or an institution. (b) Averaged over a window longer than the system’s own repair time, a positive margin keeps the trajectory inside the region where the system is still itself; a margin that stays negative for longer than that carries it out, and it stops being a member of the class. The lower band asks something different and only loosely correlated: how much of the four-pillar runtime the surviving system actually runs. That score is built so that a system missing one pillar cannot compensate with excellence in another, because that is what incompleteness means — and its weights are not derived, so the honest use of it is to rank two systems and not to certify one. The persistence class is defined against the survival of a particular packet and carries no suggestion of intention or of moral standing.

56 Why persistence and completeness belong after the full Qi stack

After Parts III–VI, the manuscript finally possesses the full Qi stack: current locking, trajectory selection, multiscale diagnostics, and memory architecture. Only at that point does it become meaningful to ask two synthesis questions. First, which open systems truly preserve themselves as continuing packets rather than as transient driven patterns? Second, among those persistent systems, how fully is the four-pillar runtime actually realized?

These are distinct questions. Persistence concerns continued viability under throughput, repair, and perturbation. Completeness concerns the degree to which the full Qi grammar is present. The revised architecture therefore places them together in Part IX. The section first isolates the persistence class SeOS, then introduces QCI as a diagnostic of realized completeness.

56.1 Historical background (with citations)

Antecedents: Schrödinger (1944), Bertalanffy (1950), Maturana–Varela (1980), Kauffman (1993), Prigogine (1967), Wilson (1975) (Maturana and Varela, 1980, Prigogine, 1967, Schrödinger, 1944); SeOS/QCI unifies persistence/completeness as the joint criterion attached to the full Qi-4 stack (see A1 Sec. 8 and A4 Sec. 7 for full history).

57 SeOS: the class of persistent self-maintaining open systems

Once the packet and ledger are defined, one can finally ask which open systems genuinely preserve themselves. Every system in nature is open, but not every open system is a self-maintaining system. A Brownian tracer in a fluid is open, but it does not protect a boundary, repair its own packet, or preserve a memory-bearing attractor. A Bénard cell is closer to self-maintenance because it exists only under continued throughput, but its packet is comparatively thin and its memory depth is limited. A cell, organism, brain, or institution is more demanding still: it must preserve not only energy throughput but also organization, boundary conditions, and records.

The purpose of SeOS is to isolate this persistent class precisely. The term “selfish” is meant in the same stripped-down technical sense in which the selfish-gene metaphor generalized local fitness logic: the system acts to maintain its own packet, not because it possesses human intention, but because the lawful runtime selects trajectories that preserve local viability whenever such preservation is possible. Cooperation, symbiosis, and hierarchy can still emerge, but they emerge as higher-order ways of stabilizing packet maintenance rather than as violations of the local logic.

To make this precise, we first need a metric notion of packet preservation.

Definition 57.1 (Packet space and viability tube). Let $\mathfrak{P}$ denote the space of admissible packets together with their current state distributions. Around a target packet-attractor $\Xi_* \in \mathfrak{P}$ define the packet distance

$$\begin{aligned} d_{\text{SiP}}^2(\Xi_1, \Xi_2) := & d_g^2(x_1, x_2) + \alpha_P \|\mathcal{P}_1 - \mathcal{P}_2\|^2 + \alpha_L \|\mathcal{L}_1 - \mathcal{L}_2\|^2 \\ & + \alpha_F |\mathcal{F}_{\text{GFF},1} - \mathcal{F}_{\text{GFF},2}|^2 + \alpha_O d_O^2(\mathcal{O}_1, \mathcal{O}_2), \end{aligned} \quad (57.1)$$

with nonnegative weights $\alpha_{P,L,F,O}$. A *viability tube* of radius ϵ is the subset

$$\mathcal{U}_\epsilon(\Xi_*) = \{\Xi \in \mathfrak{P} : d_{\text{SIP}}(\Xi, \Xi_*) \leq \epsilon, B_\partial(\Xi) \leq B_{\max}, I_m(\Xi) \geq I_{\min}\}, \quad (57.2)$$

where B_∂ measures boundary damage and I_m measures memory integrity.

Definition 57.1 says what it means for a packet to survive: the state, ports, lawful generator, ledger, and observables must remain sufficiently close to a maintained attractor, while boundary damage stays bounded and memory integrity remains above threshold.

We can now define the persistence class.

Definition 57.2 (SeOS). An open system instance with packet Ξ_t is a member of SeOS if there exist a maintained attractor Ξ_* , a viability tube $\mathcal{U}_\epsilon(\Xi_*)$, and a physically admissible feedback law such that:

- (i) **packet retention:** trajectories re-enter and remain in $\mathcal{U}_\epsilon(\Xi_*)$ after sufficiently small perturbations;
- (ii) **positive maintenance margin:** the cycle-averaged maintenance power

$$\Pi_{\text{maint}} := \langle \mathcal{P}_{\text{port}} \rangle_\tau - \langle \mathcal{D}_{\text{GFF}} \rangle_\tau - \langle \mathcal{R}_{\text{repair}}^{\min} \rangle_\tau \quad (57.3)$$

is positive for some physically relevant window τ ;

- (iii) **self-maintenance:** the repair term $\mathcal{R}_{\text{repair}}^{\min}$ includes the minimal ledger cost needed to restore packet damage, boundary mismatch, and memory degradation;
- (iv) **operational identity:** the observable dictionary remains valid while the system is maintained.

Definition 57.2 is intentionally local and non-moral. A system is “selfish” only in the sense that its lawful feedback and port usage are organized around the preservation of its own packet. It may cooperate with larger systems, embed within a larger hierarchy, or even sacrifice one subsystem to preserve a higher-order packet. What matters is that the preserved object is the system’s own effective instance packet at the chosen level of description.

58 The persistence theorem

The definition above becomes practically useful only when one can test it through a dynamical criterion. Because Part IX places persistence after the full Qi stack, that criterion must be written as a consequence of the already-built runtime rather than as an independent biological intuition. The criterion is simplest when written as a Lyapunov-type estimate.

Theorem 58.1 (Persistence criterion for SeOS(T-Qi4-14.3)). *Let Ξ_t be an admissible packet with target attractor Ξ_* . Define the viability functional*

$$\mathcal{V}(\Xi_t) := d_{\text{SIP}}^2(\Xi_t, \Xi_*) + \lambda_\partial B_\partial(\Xi_t) + \lambda_m (I_{\min} - I_m(\Xi_t))_+, \quad (58.1)$$

with $\lambda_\partial, \lambda_m > 0$. Suppose there exist constants $\kappa > 0$ and $c \geq 0$ such that along the packet dynamics,

$$\frac{d}{dt}\mathcal{V} \leq -\kappa \mathcal{V} + c - \beta \Pi_{\text{maint}}, \quad \beta > 0. \quad (58.2)$$

If the averaged maintenance margin satisfies $\Pi_{\text{maint}} > c/\beta$, then $\mathcal{V}$ remains bounded and the packet is persistent; hence the system belongs to **SeOS**. If $\Pi_{\text{maint}} < 0$ for windows longer than the intrinsic repair time, the packet leaves the viability tube and the system exits **SeOS**.

Proof. Equation (58.2) is a standard dissipative estimate. If $\Pi_{\text{maint}} > c/\beta$, then the right-hand side is strictly negative whenever $\mathcal{V}$ exceeds a finite threshold. By Grönwall's inequality, $\mathcal{V}(t)$ remains bounded and eventually enters an invariant ball around the maintained attractor. Since $\mathcal{V}$ contains packet distance, boundary damage, and memory integrity, the packet remains in the viability tube and the system is persistent. Conversely, if $\Pi_{\text{maint}} < 0$ over windows longer than the repair time, then the right-hand side eventually becomes positive, $\mathcal{V}$ grows, and either packet distance, boundary damage, or memory loss drives the system out of the tube. It therefore exits **SeOS**. ■

The theorem should be interpreted carefully. It does *not* say that every **SeOS** must be in thermodynamic equilibrium or that every member has a fixed-point attractor. The maintained attractor may be a limit cycle, torus, metastable branch, or distributed memory manifold. The theorem only says that packet preservation requires positive maintenance margin after dissipation and repair are paid for.

Table 38: Open systems, packet-bearing systems, and the **SeOS** class. Five representative systems run down the first column — a passive tracer, a driven dissipative pattern, a protocell or microbe, a neural organism, and a distributed institution — with whether each maintains a packet, its memory depth, its maintenance margin, and the resulting classification. The distinction the table draws is the sharp one: every row is open, but only those with a positive maintenance margin are self-owning, and that margin is a computable quantity rather than a judgment.

System type	Maintained packet?	Memory depth	Maintenance margin	Classification
Passive tracer / Brownian particle	No	Negligible	Not defined as self-repair	Open system, not SeOS
Driven dissipative pattern (e.g. simple convection cell)	Weak or pattern-level only	Short	Positive only while driven; thin repair logic	Borderline or proto- SeOS
Protocell / microbe	Yes	Finite adaptive memory	Positive under nutrient throughput and repair	SeOS
Neural organism / cognitive circuit	Yes	Deep internal and external memory	Positive under metabolic and information throughput	SeOS with rich Qi-HAL
Distributed institution / society	Yes, distributed	Externalized archive plus internal memory	Positive if coordination and repair exceed dissipation/shock costs	Higher-order distributed SeOS

Remark on “selfish” (non-moral). The term is non-moral: it means only that the local runtime is evaluated against preservation of the packet under study. Cooperation, symbiosis, and division of labor are therefore typical SeOS realizations; conversely, scale-dependent pathology (cancer clones, addiction loops, brittle institutions) arises whenever a sub-packet maintains its own attractor at the expense of a larger packet.

59 QCI: the Qi-4 Completeness Index

Persistence is necessary but not sufficient for the full Paper IV-B program. A packet may survive while realizing only a thin subset of the Qi runtime. A bacterium and a cortex can both be members of SeOS, but they do not instantiate the Qi pillars to the same depth. For that reason the manuscript now introduces a provisional completeness diagnostic: the *Qi-4 Completeness Index*, abbreviated QCI.

The intuition is straightforward. The four Qi pillars are not four unrelated ornaments. They are four mutually constraining layers of one runtime: current locking (Qi-QUA), trajectory selection (Qi-EP), scale diagnostics (Qi-CON), and memory architecture (Qi-HAL). A system that lacks one of these layers may still persist, but it does not realize the whole Qi-4 stack. A completeness index should therefore reward simultaneous realization rather than isolated excellence in one pillar alone.

Definition 59.1 (Qi-4 Completeness Index). For a packet-bearing open system S , let

$$q_{\text{QUA}}(S), q_{\text{EP}}(S), q_{\text{CON}}(S), q_{\text{HAL}}(S) \in [0, 1]$$

be domain-specific realization scores for the four pillars. Fix positive weights w_a with

$$\sum_{a \in \{\text{QUA}, \text{EP}, \text{CON}, \text{HAL}\}} w_a = 1.$$

The *strict* Qi-4 Completeness Index is defined by the weighted geometric mean

$$\text{QCI}(S) := q_{\text{QUA}}(S)^{w_{\text{QUA}}} q_{\text{EP}}(S)^{w_{\text{EP}}} q_{\text{CON}}(S)^{w_{\text{CON}}} q_{\text{HAL}}(S)^{w_{\text{HAL}}}. \quad (59.1)$$

The geometric mean is chosen for a reason. If one pillar is absent, completeness should collapse strongly rather than be hidden by large scores in unrelated pillars. This is exactly what the product form guarantees. If desired for phenomenology or exploratory data analysis, one may also track the additive companion score

$$\text{QCI}_+(S) := \sum_a w_a q_a(S), \quad (59.2)$$

but Eq. (59.1) is the stricter and conceptually cleaner quantity.

The pillar scores themselves are intentionally domain-dependent but structurally constrained. A useful canonical choice is:

$$q_{\text{QUA}}(S) := \exp \left[-\frac{D_{\text{lock}}(S)}{\sigma_{\text{QUA}}} \right], \quad (59.3)$$

$$q_{\text{EP}}(S) := \exp\left[-\frac{D_{\text{sel}}(S)}{\sigma_{\text{EP}}}\right], \quad (59.4)$$

$$q_{\text{CON}}(S) := \frac{1}{4}\left(c_{\text{scale}} + c_{\text{Levy}} + c_{\text{frac}} + c_{1/f}\right), \quad (59.5)$$

$$q_{\text{HAL}}(S) := \frac{1}{5}\left(c_{\text{MePMoS}} + c_{\text{NHT}} + c_{\text{HAL}} + c_{\text{BFFT}} + c_{\text{switch}}\right), \quad (59.6)$$

where the c -scores lie in $[0, 1]$, D_{lock} measures deviation from the locked-current ray, and D_{sel} measures deviation from the common selector family. Other domain-specific realizations are permitted provided they preserve the same monotonic meaning.

Table 39: Interpretation of the four pillar scores entering the completeness index. Each row is one score, with what it measures and the typical intuition for it across domains. The index is a weighted sum of the four and is offered as a comparative instrument rather than an absolute one: the paper is explicit that the weights are not derived, so the honest use of the index is to rank two systems, not to certify one.

Pillar score	What it measures	Typical intuition across domains
q_{QUA}	fidelity of current locking	whether event-level energy, matter/charge, and information bookkeeping stays near one common locked ledger
q_{EP}	selector coherence	whether observed trajectories can be organized by one admissible selector family rather than by unrelated heuristics
q_{CON}	multiscale diagnostic closure	whether scale covariance, heavy tails, fractality, and two-knee spectra co-occur coherently
q_{HAL}	architectural completeness	whether memory, sensing, motion, multiscale coding, and commitment switching form a stable holographic architecture

Remark 59.2 (Persistence is not the same as completeness). A system may be high in SeOS membership and low in QCI. A stable chemical oscillator, for example, may persist while realizing only a thin memory architecture. Conversely, a rich but damaged neural system may momentarily preserve substantial Qi structure while approaching loss of viability. SeOS and QCI therefore answer different questions: the first asks whether a packet can maintain itself; the second asks how fully the four-pillar runtime is realized on that packet.

60 Integrated view: from runtime constitution to persistence and completeness

The notions derived above and the new completeness index can now be assembled into one compact flow. An admissible system first declares a canonical packet. That packet determines the state manifold, ports, generator, ledger bundle, and observables. The generator acts on the packet distribution, while the scalar GFF functional organizes gradient descent and maintenance balance. The current vector is locked by the Arisaka Equation. The resulting dynamics then passes through

the general runtime pipeline:

$$\text{SIP} \xrightarrow{\text{declare system instance}} \text{GFF} \xrightarrow{\text{lawful GDOS evolution}} \text{DGI} \longrightarrow \text{GSSB} \longrightarrow \text{GA}. \quad (60.1)$$

If, in addition, the cycle-averaged maintenance margin remains positive, the maintained attractor is not merely a transient branch but a persistent packet:

$$(\text{SIP}, \text{GFF}, \text{GA}) + \langle \Pi_{\text{maint}} \rangle_{\tau} > 0 \implies \text{system instance } S \in \text{SeOS}. \quad (60.2)$$

Equations (60.1) and (60.2) summarize the persistence half of the present section. The completeness half is measured by QCI, which records how fully the four Qi pillars are realized on a persistent packet. Together, persistence and completeness distinguish merely surviving packets from those that realize the full memory-bearing open-system stack. The framework is no longer only a constitutional story about what theories are admissible, nor only a dynamical story about how open systems flow, nor only an architectural story about memory and prediction. It is now also a packaging story: any open system can be declared, evolved, budgeted, and tested in one grammar.

Table 40: Integrated roles of the packet, the free-energy ledger and the persistence class inside the full runtime. Four layers run down the first column, with the primary object each carries, the main mathematical task it performs, and the scientific question it answers. The table is the architecture of the paper’s second half in one page, and the fourth column is the reason the layers are in this order: each question presupposes the answer to the one above it.

Layer	Primary object	Main mathematical task	Main scientific question
Packet	SIP	Declare the admissible system instance	What exactly is the object being modeled and measured?
Ledger	GFF	Supply a unique scalar budget and its current bundle	What quantity is minimized, maintained, or spent during dynamics?
Runtime	DGI+GSSB+GA plus Qi-4	Evolve, select, diagnose, and architect the state	How does the system explore, choose, stabilize, and remember?
Persistence class	SeOS	Test invariant maintenance under throughput and repair	Does the packet merely fluctuate, or can it preserve itself as a continuing unit?

Worked examples. Persistence and completeness across protocells, organisms, brains, and institutions are tabulated with the appropriate q_{QUA} , q_{EP} , q_{CON} , q_{HAL} entries; see Appendix I and the cross-domain cards therein.

61 Staying alive, and how much of it you are doing

Two questions have been separated in this Part that are usually asked as one, and the separation is worth keeping.

The first is whether a thing can hold itself together. That turns out to be an arithmetic question. Something comes in; something is lost that cannot be recovered; something must be spent putting right what has gone wrong. If the first exceeds the other two on average, over a window long enough to include the repairs, the thing persists. If it does not, for long enough, the thing stops. There is no vital principle in that sentence and no threshold of complexity — only a margin, which can be positive or negative, and which can in principle be measured for a protocell, a coral reef, a company or a state.

The second question is how much of the machinery of the preceding Parts a surviving thing actually runs, and the answer is not correlated with the first in the way one might expect. Plenty of things persist for a very long time while doing almost nothing: a chemical clock can outlast the animal watching it. Plenty of things run a great deal of machinery while sliding toward the edge. Survival and richness are different axes, and confusing them is how the word "complexity" came to mean nothing in particular.

Both measures are deliberately narrow, and the narrowness is the honest part. The persistence class is defined against the survival of a particular packet, and says nothing whatever about whether that packet deserves to survive — a distinction that matters, because the same arithmetic describes a healing wound and a tumor, a functioning institution and a captured one. What separates those is which packet you chose to draw the boundary around, and the framework is explicit that it does not choose for you. The completeness index is narrower still: its weights are not derived, so it can rank two systems and cannot certify one.

It is worth noticing what has quietly failed to appear in this Part, because the absence is the result. There is no moment anywhere in it at which something becomes alive. There is a margin, which is positive or negative; there is an index, which is higher or lower; both vary continuously and neither has a step in it. Anyone who has tried to say where the line falls — around a virus, a prion, a self-replicating crystal, an organelle that used to be somebody else — has found that the line does not want to be drawn, and has usually concluded that the difficulty is temporary. The answer here is that the line is not there to be found. What is there is a plane with things scattered across it, and the words we have been using are names for regions of that plane rather than for kinds of thing. That is the loss of a certain sort of clarity and the gain of a more useful one, and it is the same trade this framework has made at every level: a boundary that felt absolute turns out to be a place where a continuous quantity happens to be changing quickly.

What is worth taking from this Part is smaller than a theory of life and more useful than a metaphor. It is that "staying alive" can be written as an inequality with named terms, that the terms can be measured in domains that share no vocabulary, and that whether a thing survives and whether it is doing anything interesting are two questions with two answers. Most of what makes those questions hard is that they are usually asked together.

One thing this Part has deliberately not done is look inside. The packet has been treated from without, the way a physician treats a patient or an auditor a firm: what comes in, what leaks away, what the repairs cost. But a system running the full machinery of the preceding Parts is doing something no account from outside can see. It is holding a picture of what surrounds it and correcting that picture against the evidence that reaches its own surface. The next Part turns the same open-system logic inward and asks what such a picture is, how much of the world can ever get into it, and what follows from the fact that the thing holding the picture is itself inside the world the picture is of.

Part VIII

The Qi-4 pillars as the GGA characterization

What the ground attractor carries, and in which direction it is proved

This is the shortest Part in the paper and the one that does the most joining. Two halves of a program have been built out of different materials, and this is where they are shown to be describing the same thing.

The geometric companion answers a question about shape: given a landscape, where does a system that is rolling downhill finally stop? Its answer is the deepest basin available on the stratum, and that answer is reached using nothing but the shape of the landscape. No informational condition enters it anywhere.

The four preceding Parts answer a different set of questions, about ledgers, budgets, scales and memory. They were built without reference to where anything comes to rest.

What is established here is that the configuration the geometry stops at is also the configuration that satisfies all four of the informational conditions. It carries the locked four-way ledger. It is where the unified variational principle is stationary. It is coherent across the scales between its own two turns. And it is a place where an architecture can halt rather than having to keep working. That the two descriptions, made of different materials, pick out one configuration is a real claim and it is this paper's to make.

The direction of that claim is the thing to read carefully, and the Part is deliberate about it. What is proved is that the resting place carries the four certificates. The converse — that no other basin carries all four, so that the four conditions could have found the resting place by themselves — would need an argument showing at least one condition fails at every basin that is not the deepest. On a stratum the geometric companion proves admissible that argument exists and is short: there is exactly one critical point with no direction of descent, so every other critical point fails the halting condition, and the geometry has supplied the argument rather than the four conditions. On a stratum whose admissibility is open no such argument exists in this program, and the converse is open there — the openness inherited from admissibility. It is recorded as an open problem with that scope. Asserting the converse beyond it by pointing at a companion whose own argument invokes these same four conditions would be a circle; the claim is therefore kept separated into the direction that is proved and the direction that is not.

62 The characterization theorem, pillar by pillar

The four pillars are complete, and one consolidating statement is now both possible and owed — though not the one the series has been quoting.

The Ground Geometric Attractor is not selected by the pillars. It is fixed by the landscape alone: GG-A6-GRIP's per-stratum ground-state corollary produces it from the Morse structure of

an admissible stratum, using no admissibility filter of any kind, and GG-A6-GRIP’s definition is accordingly the plain global minimum. So there is nothing left for a filter to select from once the descent has run, and a formulation that speaks of the attractor “selected from among the GA’s by the four Qi-4 pillars” misdescribes what the two papers jointly establish.

What the pillars do is *characterize* that object. The configuration the descent halts at is also the configuration that carries the locked four-current ledger, extremizes the unified variational principle, transports coherently across scales, and is one at which an architecture can halt. That the geometric description and the informational description pick out the same configuration is a substantive claim, it is this paper’s to make, and this section is its home. We state it as a theorem and make the characterization *pillar by pillar* — and we are explicit about which direction of it is proved.

Theorem 62.1 (Qi-4 characterization of the GGA (T-Qi4-GGA)). *Let Σ be an admissible stratum in the sense of GG-A6-GRIP — compact, connected, Kähler, with $\mathcal{F}_{\text{GFF}}|_{\Sigma}$ the moment map of a Hamiltonian torus action by isometries with isolated fixed points — and let $\text{GGA}(\Sigma)$ be its Ground Geometric Attractor: the unique global minimum of $\mathcal{F}_{\text{GFF}}|_{\Sigma}$, supplied by that paper’s per-stratum ground-state corollary, unique because on such a stratum every Morse index is even (Frankel, 1959), so the function is perfect and a connected stratum has exactly one critical point of index 0 (Atiyah, 1982). Then $\text{GGA}(\Sigma)$ satisfies all four Qi-4 admissibility conditions. The four act as follows:*

- **Qi-QUA (ledger admissibility).** *The candidate must carry the locked four-current identity $J_E/Q_B = J_e = J_B = J_{Q_i}$ (the Arisaka Equation): an attractor whose currents do not lock is not a lawful occupant of the boundary ledger, however stable its basin.*
- **Qi-EP (variational admissibility).** *The candidate must be a genuine extremum of the unified extremal principle — the boundary projection of the 6D least-action structure — not merely a local minimum of a coarse-grained effective potential.*
- **Qi-CON (scale coherence).** *The candidate must survive scale transport: its conformal/Lévy/ $1/f$ signatures must be those of a fixed point, so that the attractor seen at one scale is the same attractor seen at another.*
- **Qi-HAL (halting admissibility).** *The runtime must be able to halt there: an architecture descending the landscape must terminate at the candidate with no admissible descent remaining. This is the architectural face of the no-exit characterization of ground-state-ness — the GGA is where GRIP-In halts — and it is the certificate that excludes basins that cannot be reached or cannot be held: reached, because the basin’s stable set has full measure; held, because at the layer’s bath temperature its escape rate, or on an admissible stratum its equilibrium depletion, lies below the observation window (Remark 62.2).*

Existence and uniqueness of $\text{GGA}(\Sigma)$ are GG-A6-GRIP’s, not this paper’s, and are used here as an import.

Proof. Each clause is an application of the corresponding pillar theorem of Parts III–VI to the configuration GG-A6-GRIP delivers. Qi-QUA: the current locking of the Arisaka equation holds at every admissible boundary event, and the terminus of a completed descent is such an event, so

the ledger is carried there. **Qi-EP**: the terminus is a critical point of $\mathcal{F}_{\text{GFF}}|_{\Sigma}$ by construction, and by the variational theorem of Part IV the boundary projection of the 6D extremal principle is stationary exactly at such points. **Qi-CON**: by the two-knee spectral theorem the scale diagnostics of a configuration at which the descent has halted are those of a fixed point, since no residual flow remains to break scale covariance. **Qi-HAL**: the architectural halting condition asks that a descending runtime terminate with no admissible descent remaining, which is precisely what it means for the exit-channel set of the global minimum to be empty. Each of the four is therefore a property the ground attractor possesses. ■

Remark 62.2 (The four certificates in the companion’s terms). Three of the four clauses are written above in the vocabulary of the pillars, and each has an exact counterpart in the vocabulary of the descent. Stating the counterparts makes the certificates checkable where they were descriptive.

Qi-HAL. “Nothing can reach” and “nothing can stay” are two different quantities. The first is the stable set of the critical point under the descent — on an admissible stratum, the cell of [Białynicki-Birula \(1973\)](#) attached to it. **GG-A6-GRIP**’s corollary gives the floor’s cell full measure (its complement is a closed set of zero volume and empty interior) and gives every other critical point a lower-dimensional cell, so a runtime that starts anywhere generic reaches the floor and reaches nothing else. The second is what happens once there. On an admissible stratum there is no escape rate to speak of: the descent has one asymptotically stable point, no index-1 saddle and no Markov chain of metastable states, so no Eyring–Kramers formula applies, and what boundary noise at the layer’s bath temperature does is deplete the floor’s equilibrium share by an amount linear in that temperature over the height gap to the next critical point (**GG-A6-GRIP**, **GRIP-ST7**). On a stratum with several attractors — which is not admissible — the quantities are the Freidlin–Wentzell rates in and out ([Freidlin and Wentzell, 1998](#)). The certificate is therefore: a basin is halting-admissible if and only if its cell is generic and its escape rate, or on an admissible stratum its equilibrium depletion, at the layer’s bath temperature is below the observation window. The temperature is the layer’s own. **GG-A6-GRIP**’s $T_{\text{GFF}} \sim M_{\text{KK}}$ is a freeze-out record of a dissipatively disconnected block and sets nothing in a present-day layer; a cell and a brain are read at their own bath temperatures, which is what makes the certificate the same statement on every floor of Part X.

Qi-EP. The clause distinguishes a genuine extremum of the unified principle from a local minimum of a coarse-grained effective potential, and the distinction has a test. The landscape whose stationarity the certificate checks is the relative entropy to the layer’s stationary state, which [Lindblad \(1975\)](#) and [Spohn \(1978\)](#) make a Lyapunov function of any CPTP semigroup that has one, and which is $\mathcal{F}_{\text{GFF}}$, up to a constant and the temperature, exactly when that state is the Gibbs state of **GG-A6-GRIP**’s Theorem T3.2 — on an admissible stratum, the dilatation charge of that paper. A coarse-grained effective potential can be a Lyapunov function of the same flow without being a *complete* one in Conley’s sense ([Conley, 1978](#)): it can take one value on two chain components the relative entropy separates, and a local minimum of the merged landscape need not be stationary for the flow at all. Completeness is the test, and it is what the word “genuine” in the clause means.

Qi-CON. The fixed-point language of the clause is the language of continuation. A configuration whose signatures “are those of a fixed point” is one whose Conley index — the local topological data of the isolated invariant set — is invariant under the scale deformation ([Conley, 1978](#)); that is

the exact content of “the attractor seen at one scale is the same attractor seen at another,” and it is what the two-knee theorem of Part V tests. The measured record against which the signatures are audited is the subject of the second supplement ([Arisaka, 2026, S2-Conformal](#)).

Remark 62.3 (What is proved here, and what is not). The theorem proves one direction: the object GG-A6-GRIP’s descent halts at is admissible under all four pillars. The converse — that *no other* Geometric Attractor of the stratum passes all four, so that the four-pillar condition would by itself single out the ground attractor — has a scope, and the scope is inherited from GG-A6-GRIP. On an admissible stratum the converse is a theorem, and it is worth being plain about how little it costs: the corollary imported in the statement gives exactly one critical point of index 0, so every other critical point of the stratum has a direction of descent, its cell is lower-dimensional, and it fails the halting certificate of Remark 62.2. The other three certificates do no discriminating work there — every critical point is stationary, carries the ledger, and has fixed-point signatures — and the converse is the uniqueness of the floor read through Qi-HAL, at GG-A6-GRIP’s grade and with its clauses. On a stratum whose admissibility is open, the Wilson-line stratum first among those GG-A6-GRIP names, the landscape may carry two floors of equal depth, each with a generic cell and an empty exit channel; Qi-HAL then cannot tell them apart, whether the other three can is exactly what is not known, and the converse is open. The openness is inherited from admissibility, not produced by the pillars; it is recorded as **PS-Qi4-10**, with that scope. (**PS-Qi4-6**, to which this gap was once assigned, is the uniqueness of the Qi-HAL factorization and a different question.) Earlier formulations of this section asserted the converse without scope and rested it on the capstone, whose own proof of the identification invokes these pillars — a circle, and the reason the claim is now separated into a proved direction and an open one. Nothing downstream is weakened by the separation, because every downstream use is of the proved direction: papers quote the four certificates as properties of the electron, the proton and the Kaxion, never as the instrument that found them.

Physical picture. Chemistry again supplies the intuition. A molecule has many stationary electronic configurations; the ground state is the one that carries the right electron count (ledger), genuinely minimizes the energy (variational), looks the same under every spectroscopic magnification (scale coherence), and is where relaxation actually ends (halting). The Qi-4 pillars are those four certificates promoted to constitutional rank. The worked case is the electron. Its stratum is the charged-lepton stratum, which GG-A6-GRIP establishes as admissible, so the ground attractor exists and is unique there by that paper’s corollary; the electron is its readout; and the theorem above says the four certificates are all carried at that configuration — the locked ledger, the stationarity, the fixed-point diagnostics, and the empty exit channel that is the reason it does not decay. That is why every electron, every proton and every Kaxion is the same object everywhere: not because a filter picked one out of many, but because there is one ground attractor per stratum and the four certificates describe it ([Arisaka, 2026, A9-SM](#); [Arisaka, 2026, C1-Kaxion](#)).




Figure 14: Where the two halves of the program meet, and — as importantly — which way the meeting runs. The geometric companion fixes the resting place using the shape of the landscape alone; no informational condition enters that argument anywhere. What this Part establishes is that the configuration the descent halts at is also the configuration satisfying all four of the informational conditions built in the preceding Parts: it carries the locked four-way ledger, it is where the single budget is stationary, it is coherent across the scales between its own two knees, and it is a place where an architecture can stop rather than having to keep working. That two descriptions made of entirely different materials pick out one configuration is a substantive claim and it is this paper's to make. The dashed return arrow is the direction that is not established on its own: whether the four conditions could by themselves have found that configuration, with the geometry set aside. Where the companion proves the stratum admissible the answer is yes, and the geometry supplies it — there is one critical point with no way down, so every other fails the halting condition; where admissibility is open the answer is open, no argument exists in this program, and it is recorded as an open problem with that scope rather than quoted as settled.

63 Four certificates, and what they are not

A chemist already knows what this Part has done, in a smaller setting. Ask why a molecule sits in the state it sits in, and the answer has several parts that arrive together. It carries the right number of electrons. It genuinely minimizes the energy rather than merely sitting at a flat spot. It looks like itself under every instrument that can be pointed at it. And it is where relaxation actually stops, rather than a way station on the road to somewhere lower. Four certificates, held at once, by one configuration.

The four pillars are those certificates promoted to constitutional rank, and made to apply to anything that persists rather than only to molecules. The ledger is Nature's bookkeeping in four currencies at once. The variational principle is the statement that the path taken, the currents maintained, and the beliefs updated are three readings of one budget. The scale condition is what the thing must look like between its slowest and its fastest surviving rate. And the architecture is what a system must become once it holds a memory for longer than one act of looking.

It matters what kind of claim this is, and the Part is deliberate about it. The certificates describe the resting place; they did not select it. The descent finds where to stop by the shape of the landscape alone, with no informational filter applied anywhere. What is established here is that the place it stops at satisfies all four conditions — and that is a real claim, since the geometric description and the informational description are built out of entirely different materials, and there is no obvious reason they should agree on a single configuration. Whether the four conditions could by themselves have found that configuration, with the geometry set aside, is established exactly where the companion proves the landscape has one floor, and there it is the geometry that establishes it — the one place with no way down is the only place an architecture can halt. Where the companion has not proved that, it is not established. It is written down as an open problem with that scope, and it should not be quoted as settled beyond it.

That distinction is worth more than it costs. It is the reason every electron is the same as every other electron, and every proton the same as every other proton — not because some filter selected one candidate out of many, but because there is exactly one place to stop, and these four statements are what a reader would find written on it.

There is a smaller lesson in the constant. Information has spent a century as the one quantity in science with no price attached: freely invoked, freely conserved or destroyed as the argument required, and never billed. Here it is billed, at a rate fixed by the same integers that fix how many families of matter there are, and the bill can be checked against a molecule that every living thing on this planet carries. Whether the check passes is a matter for the laboratory. That it is a check at all — that a statement about bits can be settled by a calorimeter — is the change this Part is arguing for.

Part IX

Internal Measurement and the EIC Runtime

Outer Hologram, Inner Hologram, and the emergence of an internal cosmos

Nothing ever touches the world. It touches its own boundary.

That is not a philosophical flourish; it is what the projection of the preceding Parts already implies, and this Part follows the implication. A retina receives a two-dimensional sheet of light and the animal behind it needs three dimensions to move through. A bacterium samples one concentration at one point and needs to know which way the food is. A physicist reads four-dimensional instruments and is trying to establish something about six dimensions. In every case the thing that guides the next action cannot have been handed over intact. It has to be built.


So this Part separates three objects that are usually run together. The first is what the world lawfully delivers to a boundary — real, public, and lossy, since many different external situations deliver the identical record. The second is what a system maintains internally: not a copy of its surroundings but a working model with a geometry and a memory, which is a very different sort of thing from a photograph. The third is the loop that keeps the second answerable to the first.

Two results here are worth stating in advance because they overturn a natural assumption. The first is that motion is not an accessory to perception but a condition of it: a system that cannot act on its surroundings cannot, in general, recover them at all, since different external states would produce identical sensory histories, and moving is what breaks the tie. A retina that cannot move is depth-poor; a bacterium that cannot swim is gradient-poor. The second is that memory is not optional either — without it there is no internal world, only a running estimate that forgets itself.

The Part then does something it would be easy to do badly, and takes some care not to. It defines a technical term for the portion of a system's internal model that has been selected, stabilized, and made available to the system's own next cycle. The definition is operational, meant to be applied to data, and graded rather than all-or-nothing. It is not a theory of subjective experience, it takes no position in the debates about one, and the table of non-claims in this Part is not a disclaimer bolted onto the results — it is what makes the results sayable at all. Most of what this Part contributes is organization; the short table of predictions marks the few rows that the world could actually contradict.

64 OH+ IH= EIC: Outer Hologram, Inner Hologram, and Internal Measurement

With persistence and completeness fixed, the theory can turn to its deepest conceptual distinction: external projection versus internal reconstruction. Section 64 therefore introduces OH, IH, and EIC as the final conceptual layer of Part IX. This placement is deliberate. Internality must come after



Action-conditioned observability theorem (T-HAL): when action is restricted, IH model quality degrades.
Memory-bearing $U(1)_{Q_i}$ systems realize the closed loop; SeOS systems persist within it.

Figure 15: The Outer Hologram + Inner Hologram = Evolving Inner Cosmos ($OH + IH = EIC$) framework of Part IX. The Outer Hologram OH is the environment encoded on the system’s boundary (’t Hooft 1993, Susskind 1995); the Inner Hologram IH is the system’s internally maintained model of that environment (Llinás 2001, predictive coding). The Evolving Inner Cosmos EIC is the closed-loop coupling $C_{loop}(OH, IH)$ that the conscious organism inhabits. Per the clarification, the formula “ $OH + IH = EIC$ ” is definitional shorthand for $EIC := C_{loop}(OH, IH)$ rather than an arithmetic identity; the structural content lives in the Lyapunov refinement theorem on EIC monotonicity.

runtime constitution, Qi architecture, persistence, and completeness; otherwise it risks floating free of the lawful machinery that makes it possible.

64.1 Why the runtime still needs a distinction between external projection and internal reconstruction

After everything from Section 4 to Part VII, the manuscript already possesses a universal runtime, a current ledger, a selector principle, a diagnostic layer, a memory architecture, a packet vocabulary, and a persistence criterion. Yet one decisive conceptual task remains unfinished. The earlier sections explain how an open system evolves, how it budgets its currents, how it stores memory, and how it persists. They do not yet fully explain how the world appears *to* an open system, how an open system reconstructs that world internally, and how a selected internal reconstruction may come to function as conscious content.

This missing step is exactly what the relation

$$OH + IH = EIC$$

is designed to capture. The notation is intentionally compact, but it must not be read superficially. The symbol OH does not mean “the world itself.” It means the lawful projection of the inaccessible or higher-dimensional world into observer-accessible boundary records. The symbol IH does not mean “a copy of the world in the head.” It means the internal reconstructive manifold through which a memory-bearing open system infers, stabilizes, and reuses an allocentric model of what lies beyond its immediate sensory boundary. The symbol EIC denotes the *Evolving Inner Cosmos*: the closed loop in which externally projected evidence and internally reconstructed structure are continually aligned, updated, selected, and stabilized.

The need for this distinction is not merely philosophical. It is forced by the mathematics of partial access. A finite open system never interacts with the total six-dimensional bulk in its full state-description. It interacts only through ports and measurement channels. Likewise, a living or cognitive system never receives the full external state; it receives only a narrow sensory stream, often lower-dimensional than the structure it must navigate. The world that guides action must therefore be internally reconstructed rather than directly given. The present section formalizes that statement.

It is useful to mark the three layers immediately. We write H_t^{out} for the concrete time-dependent outer hologram, H_t^{in} for the time-dependent inner hologram, and $\mathfrak{E}_t$ for the corresponding evolving inner-cosmos state. Table 41 gives the high-level distinction.

Table 41: Basic distinction between the outer hologram, the inner hologram and their coupling. Each row is one object, with what it is, what it is *not*, and its constitutional role. The middle column is the important one and is written to close the two most common misreadings: the outer hologram is not a complete description of the environment, and the inner one is not a picture in a viewer’s head. Both are structures on a boundary, and the third object is their coupling, not their sum.

Object	What it is	What it is not	Constitutional role
OH	Lawful projection from the 6D bulk to accessible boundary records	Not the full bulk state itself	Extends OGGEF from laboratory measurement to all scientific observation
IH	Internal reconstructive manifold built from memory, prediction, action, and sensory history	Not a literal pixel-by-pixel copy of the external world	Explains how open systems build allocentric models from partial data
EIC	Closed-loop coupled process aligning OH and IH through repeated updates	Not a mystical addition to dynamics	Explains internal world-model formation, conscious content, and learning across scales

The section therefore proceeds in five steps. First, it derives the outer hologram directly from the OGGEF boundary law. Second, it derives the inner hologram from the inverse problem faced by any memory-bearing open system with action. Third, it shows why motion, memory, and allocentric stabilization are mathematically necessary if one starts from low-dimensional or egocentric sensory data. Fourth, it gives a disciplined interpretation of the relation $\text{OH} + \text{IH} = \text{EIC}$ by introducing

an explicit alignment functional. Fifth, it states the consciousness thesis in its careful technical form: consciousness is the system's internal measurement of a selected and stabilized portion of the inner hologram down to an internal boundary implemented by MePMoS and the broader Qi-HAL architecture.

64.2 The Outer Hologram: lawful projection from the 6D bulk to measurable boundary records

GG-A5-GDOS already supplies the BulkMap $\mathcal{D}_{\text{OGGEP}} : \text{bulk data} \longrightarrow (\mathcal{H}_\partial, \rho_\partial, \Phi_{t,s}, \{\mathcal{I}_\alpha\})$ demanded by AX-3; see GG-A5-GDOS Theorem T-OGGEP for the derivation. The present subsection re-reads it as the outer-hologram definition. Define the outer-hologram kernel

$$K_{\text{OH}}(y|X; t) := \text{Tr}[\mathcal{I}_y \Phi_{t,0}(|X\rangle\langle X|)] , \quad (64.1)$$

with $X \in \mathcal{X}_6$ a bulk microstate and $y \in \mathcal{Y}$ an observer-accessible boundary record. In a classical limit this reduces to an ordinary stochastic kernel; in a quantum setting it is the instrument-induced probability density of y given X . The actual outer-hologram record at time t is

$$H_t^{\text{out}}(y) = \int_{\mathcal{X}_6} K_{\text{OH}}(y|X; t) \rho_6(X, t) dX, \quad (64.2)$$

with normalization

$$K_{\text{OH}}(y|X; t) \geq 0, \quad \int_{\mathcal{Y}} K_{\text{OH}}(y|X; t) dy = 1. \quad (64.3)$$

Definition 64.1 (Outer Hologram). Given a bulk state family $\rho_6(X, t)$, reduced boundary evolution $\Phi_{t,0}$, and admissible instrument class $\{\mathcal{I}_y\}$, the *outer hologram* at time t is the boundary law H_t^{out} of Eq. (64.2); an individual registered observation y_t is an *OH event*.

The outer hologram is *lawful* (fixed by the constitutional map, not an arbitrary representation), *lossy* (many bulk states share the same boundary law), and *instrument-relative* (different instrument families coarse-grain the same bulk reality differently). This is not subjectivity but admissible boundary indexing.

Proposition 64.2 (The outer hologram is a quotient object). *With $X \sim_{\text{OH}} X'$ iff $K_{\text{OH}}(\cdot|X; t) = K_{\text{OH}}(\cdot|X'; t)$ for the chosen instrument class, the outer hologram depends only on the equivalence class $[X]_{\text{OH}}$: observation accesses a quotient of bulk reality.*

Proof. Eq. (64.2) depends on X only through $K_{\text{OH}}(y|X; t)$; equivalent bulk states give identical contributions for every admissible boundary record. ■

Every registered scientific measurement is an OH event: a particle-detector click, an image frame, an ecological count, a spike train, a fluorescence reading. Empiricism says observation is the source of knowledge; GDOS says observation is a lawful OH event already shaped by the bulk-to-boundary map, and the IH (next subsection) is what turns a stream of such events into an internally usable world.

Table 42: Examples of outer-hologram events across scientific domains. Four domains run down the first column, with the external structure being encoded, the boundary channel or instrument that does the encoding, and what constitutes an event in each. The table’s purpose is deflationary: the outer hologram is not an exotic object but the ordinary situation of a finite system learning about what it cannot directly access, and the third column is always a real instrument or a real membrane.

Domain	Underlying external structure	Boundary channel / instrument	OH event
Particle physics	Bulk field excitation or scattering process	Detector, trigger, reconstruction chain	Hit pattern, track, calorimeter energy, event class
Cell biology	Molecular state, conformational or concentration field	Receptor, microscope, fluorophore, reporter construct	Bound receptor state, image frame, fluorescence time point
Neuroscience	External world plus organism pose and context	Retina, cochlea, mechanoreceptors, electrodes	Sensory sample, spike train, local field potential record
Ecology / society	Hidden population or collective state	Survey protocol, sensor network, archive	Count, flow value, transaction record, communication event

64.3 The Inner Hologram: internal reconstruction of allocentric structure from partial sensory streams

The outer hologram delivers data, but no open system can survive or act intelligently by storing raw data alone. If the incoming stream is partial, delayed, noisy, or lower-dimensional than the relevant external structure, a reconstructive model is required. That model is the inner hologram.

Let $x_t \in \mathcal{X}_{\text{env}}$ denote the external state relevant to the system at time t , $a_t \in \mathcal{A}$ the action or motion variable, $s_t \in \mathcal{S}$ the sensory sample, and $m_t \in \mathcal{M}$ the internal memory state. We consider the generic controlled open-system dynamics

$$x_{t+1} = f(x_t, a_t) + \omega_t, \quad (64.4a)$$

$$s_t = h(x_t) + \nu_t, \quad (64.4b)$$

with process perturbation ω_t and measurement noise ν_t . In many cases h is strongly non-injective: different external states generate indistinguishable instantaneous sensory samples. This is the rule rather than the exception. A 2D retina does not uniquely determine a 3D world at one instant; a single odor concentration sample does not uniquely determine a surrounding concentration field; a scalar temperature reading does not uniquely determine a thermal landscape.

The system therefore introduces an internal latent variable $z_t \in \mathcal{Z}$ that does not merely store the last s_t but represents a structured hypothesis about the external world. A natural variational definition is

$$q_t(z_t) = \arg \min_q \left\{ \mathbb{E}_q[-\ln p(s_t|z_t)] + D_{\text{KL}}(q(z_t) \parallel p(z_t|m_{t-1}, a_{t-1})) \right\}, \quad (64.5)$$

which is the Qi-EP/FEP-type update written specifically for internal world-model reconstruction. The first term fits the current sensory evidence, while the second term preserves continuity with the previously predicted state. In plain language: the system asks which internal latent configuration best explains the incoming OH event without discarding what its past memory and action already predicted.

Definition 64.3 (Inner Hologram). For a memory-bearing open system, the *inner hologram* at time t is the structured internal reconstructive object

$$H_t^{\text{in}} := (\mathcal{Z}, q_t, g_t, W_t), \quad (64.6)$$

where $\mathcal{Z}$ is the latent internal manifold, q_t is the current posterior or belief state on $\mathcal{Z}$, g_t is the learned or reconstructed metric/geometry on that manifold, and W_t is the HAL memory substrate storing the relational attractor structure.

Definition 64.3 is intentionally broader than a probabilistic belief state. The IH is not only a posterior density. It also contains geometry and memory. The metric g_t enters because the internal model must encode distances, directions, or relational structure relevant for prediction and action. The memory tensor W_t enters because the reconstructed world must persist across time; without storage there is no genuine internal world-model, only a transient estimate.

The role of the Qi-HAL modules is now immediate. MePMoS provides the causal update order. NHT provides a phase-like coordinate system through which latent geometry can be encoded. HAL provides the storage substrate of recurrent relational structure. BFFT coordinates scales, so that fast sensory detail and slow contextual priors can coexist in the same IH. The DGI-GSSB switch governs how multiple candidate interpretations remain coexistent during exploration and then collapse to a selected interpretation once evidence and recurrent gain cross threshold.

The IH update may be written in MePMoS-compatible form as

$$\hat{z}_{t|t-1} = \mathcal{P}(m_{t-1}, a_{t-1}), \quad (64.7a)$$

$$q_t = \arg \min_q \left\{ \mathbb{E}_q[-\ln p(s_t|z_t)] + D_{\text{KL}}(q(z_t) \parallel \delta_{\hat{z}_{t|t-1}}) \right\}, \quad (64.7b)$$

$$m_t = \mathcal{U}(m_{t-1}, q_t, s_t), \quad (64.7c)$$

$$a_t = \pi(m_t, q_t), \quad (64.7d)$$

where $\delta_{\hat{z}_{t|t-1}}$ denotes the predicted latent prior concentrated around the predicted state, and π is the action policy. Equations (64.7a)–(64.7d) make clear that the IH is not a static map. It is a continually updated reconstructive manifold embedded in the action-perception loop.

Remark 64.4. The word “inner” should not mislead. The IH is not private in principle because it can leave external traces through action, report, physiology, and behavior. It is “inner” only in the operational sense that it is realized inside the system as the latent manifold guiding action, rather than in the external world directly.

64.4 Why motion and memory are mathematically necessary for the inner hologram

The most common intuitive mistake is to imagine that an internal world-model can be built from raw sensing alone. The mathematics says otherwise. Whenever the instantaneous sensory map h is non-injective, unique reconstruction from one time slice is impossible. History and controlled motion are required.

To formalize this, linearize Eqs. (64.4a)–(64.4b) around a nominal trajectory:

$$\delta x_{t+1} = A_t \delta x_t + B_t \delta a_t, \quad (64.8a)$$

$$\delta s_t = C_t \delta x_t. \quad (64.8b)$$

Over a horizon T , the action-conditioned observability matrix is

$$\mathcal{O}_T = \begin{bmatrix} C_0 \\ C_1 A_0 \\ C_2 A_1 A_0 \\ \vdots \\ C_T A_{T-1} \cdots A_0 \end{bmatrix}. \quad (64.9)$$

Theorem 64.5 (Action-conditioned observability of the IH (T-Qi4-15.5)). *Assume the linearized model (64.8) is valid on a finite horizon $0 \leq t \leq T$. If $\text{rank}(\mathcal{O}_T) = \dim(\delta x_0)$, then the external state perturbation δx_0 is uniquely determined by the sensory history $(\delta s_0, \dots, \delta s_T)$ and the action-conditioned state transition sequence (A_t, B_t) , up to any symmetry already quotiented by the observation map. If $\text{rank}(\mathcal{O}_T) < \dim(\delta x_0)$, then there exist distinct external perturbations that yield identical sensory histories, so no unique IH can be reconstructed from the available history.*

Proof. The sensory sequence generated by δx_0 is

$$\begin{bmatrix} \delta s_0 \\ \delta s_1 \\ \vdots \\ \delta s_T \end{bmatrix} = \mathcal{O}_T \delta x_0 + \text{known terms generated by the chosen action sequence.}$$

If $\mathcal{O}_T$ has full column rank, the nullspace is trivial and δx_0 is uniquely recoverable from the sensory history modulo symmetries already removed by the map. If the rank is deficient, the nullspace is nontrivial; hence there exists $\delta x_0 \neq 0$ such that $\mathcal{O}_T \delta x_0 = 0$. This nonzero perturbation changes the external state but produces no change in the sensory history, proving non-identifiability. ■

Theorem 64.5 is the rigorous reason why motion is not an accessory but a constitutive part of perception. The underlying rank condition is the classical observability criterion of Kalman (1960); the contribution here is its action-conditioned application to internal-model reconstruction, continuing the active-perception lineage of Bajcsy (1988) and O'Regan & Noë (2001). Self-motion changes the observation geometry and thereby increases the rank of the accumulated observability

matrix. A retina without motion is depth-poor. A bacterium without swimming is gradient-poor. A brain without active sensing is world-model-poor. This is precisely why MePMoS must contain motion before the next sensory registration.

Memory is equally necessary. The observability matrix acts over a history window, and history must be retained somewhere. A system with no memory cannot accumulate the multi-view evidence needed to invert a non-injective sensory map. Therefore the very existence of an IH already presupposes the packet, the ledger, and the memory architecture introduced earlier in the paper.

Corollary 64.6 (No genuine IH without history). *If the system has no persistent state capable of retaining information about prior sensory samples or prior actions, then any internal state at time t is a function only of the current sample s_t . In that case the internal representation cannot generally invert a non-injective observation map and therefore cannot constitute a genuine inner hologram except in trivial fully observable environments.*

Proof. If no persistent state exists, the internal representation is memoryless: $z_t = \varphi(s_t)$. Whenever h is non-injective, there exist $x \neq x'$ with $h(x) = h(x') = s_t$. Then $\varphi(s_t)$ assigns the same internal state to both external possibilities, so it cannot reconstruct the latent structure uniquely. Only if h is injective already, or the relevant task ignores the missing variables, can a memoryless representation suffice. ■

The corollary is especially important for the broad application of the theory to all open systems with memory. The moment a system becomes memory-bearing and action-bearing under nonzero J_Q , the mathematical door to IH formation opens. This does not guarantee a rich cognitive map. It guarantees only that an internal reconstructive manifold is now a lawful possibility rather than a category mistake.

64.5 Why the inner hologram must be allocentric rather than a stored sensory movie

The next issue is subtle but fundamental. Why is the IH not merely a compressed recording of sensory history? Why insist that it is allocentric or world-centered?

The answer is that action-guiding reconstruction must factor out self-motion. Let $g_t \in SE(n)$ denote the system's pose (position and orientation) in an n -dimensional task space, and let ξ_t denote a world-centered latent configuration. A generic egocentric observation has the form

$$s_t = \pi(g_t^{-1} \cdot \xi_t) + \nu_t, \quad (64.10)$$

where π denotes projection through the sensory channel. If the internal model stores only s_t , then every self-motion $g_t \mapsto g_{t+1}$ changes the entire representation even when the external world ξ_t is unchanged. Such a representation is unsuitable for stable prediction, navigation, or recognition across viewpoints.

The inner hologram must therefore transport old information into a pose-invariant frame. A minimal update law is

$$\hat{\xi}_{t+1} = T(a_t) \hat{\xi}_t + K_t \epsilon_{t+1}, \quad \epsilon_{t+1} = s_{t+1} - \hat{s}_{t+1|t}, \quad (64.11)$$

where $T(a_t)$ is the transport induced by self-motion, and K_t is an evidence gain. Equation (64.11) is the simplest mathematical expression of allocentric reconstruction. The old map is not discarded after motion; it is transformed by the internally known or inferred motion and then corrected by new evidence.

This is why spatially and temporally structured codes arise naturally in neural systems. Place-like states encode location in a world-centered frame, head-direction or ring-attractor variables encode orientation, and phase-structured sequences encode ordered trajectories or temporal position (O’Keefe and Dostrovsky, 1971, Turner-Evans et al., 2020). The present paper does not require one biological implementation only. It requires the mathematical role those implementations play: self-motion-compensated reconstruction of an external structure that remains stable despite changing sensory views.

Proposition 64.7 (Allocentric stability criterion). *Suppose the external latent configuration ξ_t is constant while the system changes pose g_t . A purely egocentric internal code $r_t = \varphi(s_t)$ is viewpoint-stable only if $\varphi \circ \pi(g_t^{-1} \cdot \xi)$ is independent of g_t for all relevant poses, which is generically impossible for nontrivial projection maps. An allocentric update of the form (64.11), by contrast, preserves the representation of ξ under self-motion up to bounded correction error.*

Proof. For nontrivial projection π , different poses g_t produce different egocentric observations of the same ξ . Therefore a direct code $r_t = \varphi(s_t)$ changes with pose unless φ performs the full inverse transport for all g_t , which is exactly what an allocentric reconstruction does. Under Eq. (64.11), the old estimate is first transported by the known action-induced map $T(a_t)$, leaving the encoded world configuration invariant in the absence of noise; the correction term $K_{t \leftarrow t+1}$ accounts only for residual mismatch. Hence viewpoint stability is achieved up to bounded prediction error. ■

In everyday language the distinction is simple. The IH is not a film running across the senses. It is a map. In richer systems it is a hierarchy of maps: spatial, temporal, causal, social, symbolic, and theoretical. A physicist reconstructing a six-dimensional theory from four-dimensional experiments is performing a very high-level IH operation. The same architectural logic begins, in primitive form, whenever a bacterium’s adaptive state stores something more useful than its last measurement alone.

64.6 The equality $\text{OH} + \text{IH} = \text{EIC}$: precise meaning of the plus sign

We are now in a position to interpret the equation itself. The expression

$$\text{OH} + \text{IH} = \text{EIC}$$

should not be read as scalar addition. The plus sign denotes *closed-loop lawful composition on a common operational boundary*. The outer hologram supplies evidence. The inner hologram supplies reconstruction. The evolving inner cosmos is their repeatedly updated coupling.

Definition 64.8 (Evolving Inner Cosmos). For a system with outer hologram H_t^{out} and inner hologram H_t^{in} , the *Evolving Inner Cosmos* is the coupled state

$$\mathfrak{E}_t := H_t^{\text{out}} + H_t^{\text{in}} \equiv \mathcal{C}_{\text{loop}}(H_t^{\text{out}}, H_t^{\text{in}}), \quad (64.12)$$

where $\mathcal{C}_{\text{loop}}$ denotes closed-loop composition through prediction, action, sensing, memory update, and internal readout. The notation $+$ is therefore mnemonic for coupling rather than arithmetic addition.

To make the coupling quantitative, let the OH-induced sensory law at time t be

$$p_t^{\text{OH}}(s) = \int_{\mathcal{X}_6} K_{\text{OH}}(s|X; t) \rho_6(X, t) dX, \quad (64.13)$$

and let the IH-induced predicted sensory law be

$$\hat{p}_t^{\text{IH}}(s) = \int_{\mathcal{Z}} p(s|z, a_t) q_t(z) dz. \quad (64.14)$$

Their mismatch is measured by the reconstruction divergence

$$\Delta_{\text{rec}}(t) = D_{\text{KL}}(p_t^{\text{OH}} \| \hat{p}_t^{\text{IH}}). \quad (64.15)$$

When $\Delta_{\text{rec}}(t)$ is small, the inner hologram predicts the OH well. When it is large, the internal model poorly matches the external evidence.

Because internal reconstruction is not free, one must add the runtime cost already derived earlier in the paper. A convenient composite functional is

$$\mathcal{F}_{\text{EIC}}(t) = \Delta_{\text{rec}}(t) + \lambda_E \mathcal{F}_{\text{GFF}}(t) + \lambda_C \mathcal{C}[q_t], \quad (64.16)$$

where $\mathcal{F}_{\text{GFF}}$ is the free-energy ledger of Sections 8 and 9, and of Part VII, $\mathcal{C}[q_t]$ is a complexity or model-maintenance term such as $D_{\text{KL}}(q_t \| q_{\text{prior}})$, and $\lambda_E, \lambda_C > 0$ are weights. Equation (64.16) says in the clearest possible way what the EIC does: it aligns internal reconstruction with outer evidence, but only under lawful energetic and entropic budget constraints.

If θ_t denotes the collection of internal parameters entering q_t , g_t , W_t , and the action policy, then the Qi-EP-consistent gradient normal form is

$$\dot{\theta}_t = -G_\theta^{-1} \nabla_\theta \mathcal{F}_{\text{EIC}} + \eta_t, \quad (64.17)$$

with G_θ a positive metric on parameter space and η_t a bounded fluctuation term.

Theorem 64.9 (Monotonic refinement of the EIC in the noiseless normal form (T-Qi4-15.9)). *Assume G_θ is positive definite and $\eta_t = 0$. Then the EIC functional satisfies*

$$\frac{d}{dt} \mathcal{F}_{\text{EIC}}(t) = -\nabla_\theta \mathcal{F}_{\text{EIC}}^\top G_\theta^{-1} \nabla_\theta \mathcal{F}_{\text{EIC}} \leq 0. \quad (64.18)$$

Hence $\mathcal{F}_{\text{EIC}}$ is a Lyapunov functional for the coupled OH-IH update and decreases monotonically until the system reaches a stationary point or attractor basin.

Proof. Differentiate Eq. (64.16) along the flow (64.17):

$$\frac{d}{dt}\mathcal{F}_{\text{EIC}} = \nabla_{\theta}\mathcal{F}_{\text{EIC}}^{\top}\dot{\theta}_t = -\nabla_{\theta}\mathcal{F}_{\text{EIC}}^{\top}G_{\theta}^{-1}\nabla_{\theta}\mathcal{F}_{\text{EIC}},$$

since $\eta_t = 0$. Because G_{θ}^{-1} is positive definite, the quadratic form is nonnegative, proving Eq. (64.18). ■

Theorem 64.9 is deliberately local and normal-form in spirit. It does not prove that every real system globally reconstructs the true external world in all circumstances. It proves the more careful statement actually needed for the paper: once the runtime has an OH channel, an IH manifold, and a Qi-EP-compatible update rule under GFF budget constraints, there is a canonical coupled functional whose decrease corresponds to progressive internal alignment with externally projected evidence. That progressively aligned loop is the EIC.

Remark 64.10. The EIC does not recover the full six-dimensional bulk “as it is in itself.” It recovers, at best, an operationally adequate model of the equivalence class of external states accessible through the system’s boundary channels. This is not a weakness of the theory; it is the precise consequence of Proposition 64.2.

The evolving-inner-cosmos language now acquires its disciplined meaning. The “cosmos” is the system’s internal world-model, not because it contains everything, but because it is the internal totality relative to which the system navigates, predicts, and acts. It is “evolving” because it is continually updated through the OH-IH loop. Human scientific learning is simply one very elaborate case of this general process: measurements enlarge and correct the inner world-model so that it better matches the externally projected lawful structure.

64.7 Consciousness as internal measurement of the inner hologram

A note on terminology. This subsection uses the word “consciousness” only in the precise technical sense defined below, namely as an operation on the inner hologram. It does not engage with broader interpretive debates about subjective experience, qualia, panpsychism, or higher-order theories; readers from those fields may treat the present subsection as defining a particular operational class without prejudice to those debates. The technical content is, deliberately, an operational criterion — something that can be applied to neural data, to behavioral data, and (eventually) to the engineered systems of Section 67 — and not a metaphysical claim.

The final conceptual step is the most delicate. The present framework does not identify consciousness with the OH, nor with raw sensory inflow, nor with the mere existence of an IH. Consciousness is a special operation performed *on* the IH.

The correct analog of external measurement is internal measurement. Externally, the bulk is projected to a boundary record through OGGEF. Internally, a selected region of the IH is projected to a self-accessible internal boundary where it can bias future prediction, future action, memory writing, and in rich organisms explicit report.

Let $\rho_{\text{IH}}(t)$ denote the internal state on the IH manifold. Let $\{\Pi_c\}$ be an admissible family of coarse-grained internal readout operators or projectors satisfying

$$\Pi_c \geq 0, \quad \sum_c \Pi_c = I_{\text{IH}}.$$

The internal measurement statistics are then

$$p_t(c) = \text{Tr}(\Pi_c \rho_{\text{IH}}(t)), \quad (64.19)$$

and the post-selected internal content state for channel c is

$$\rho_{\text{IH}}^{(c)}(t) = \frac{\Pi_c \rho_{\text{IH}}(t) \Pi_c}{\text{Tr}(\Pi_c \rho_{\text{IH}}(t))}. \quad (64.20)$$

Definition 64.11 (Internal measurement). A system *internally measures* its inner hologram at time t if there exists an internal readout boundary ∂_{int} and a family of readout channels $\{\Pi_c\}$ such that at least one selected internal content state $\rho_{\text{IH}}^{(c)}(t)$ re-enters the MePMoS loop as a causal variable for subsequent memory update, action selection, or report.

Definition 64.11 is the formal version of the informal description of consciousness as “our internal measurement of IH down to the 4D boundary by MePMoS.” The key idea is recursion: the IH is itself turned into a boundary-accessible object for the system. The system does not only model the world; it also samples or measures the state of its own model.

This is exactly where the earlier runtime layers reappear. During DGI exploration, several candidate interpretations of the current data may coexist in the IH. Through GSSB, one interpretation becomes selected. Through GA, the selected interpretation becomes stable enough to influence subsequent processing. A conscious content, in the present technical sense, is not the entire cloud of preselection possibilities but the selected and stabilized internal branch that is accessible on the internal boundary.

Definition 64.12 (Conscious EIC episode). A *conscious EIC episode* occurs at time t when all of the following hold:

- (a) the system possesses a nontrivial inner hologram H_t^{in} ;
- (b) at least one internal content channel c_* is selected from competing candidates by a DGI→GSSB→GA sequence;
- (c) the selected content persists on a timescale sufficient for writing, reuse, or action bias, i.e. it is not an instantaneous unstable fluctuation;
- (d) the selected content is available on an internal boundary ∂_{int} for re-entry into MePMoS as memory, action, or report bias.

Definition 64.12 is intentionally careful. It does *not* say that every internal model is conscious. It does *not* say that all information processing is conscious. It says only that conscious content is

the internally measured, selected, and stabilized part of the IH that becomes causally available to the next cycle of the system's own runtime.

The operational distinction matters. A great deal of sensory and neural processing can remain subthreshold, fragmented, or preselective. In the present language such processing lives in the IH but has not yet crossed the $\text{DGI} \rightarrow \text{GSSB} \rightarrow \text{GA}$ chain needed for internal measurement on ∂_{int} . Conversely, once such a crossing occurs and the content becomes self-available, one has a candidate conscious episode.

Proposition 64.13 (Consciousness is not raw sensation). *Under Definitions 64.3, 64.8, and 64.12, raw sensory samples s_t are neither necessary nor sufficient for conscious content. They are not sufficient because they may fail to be integrated into an IH and internally measured. They are not necessary in the strict momentary sense because internally reactivated or imagined content can be selected from stored IH structure even when current sensory drive is weak or absent.*

Proof. Insufficiency: by Definition 64.12, conscious content requires a selected and stabilized internal branch accessible on ∂_{int} , not merely the existence of a sensory sample. Necessity failure in the strict momentary sense: the same definition allows internally driven content states selected from H_t^{in} if they are available on the internal boundary for subsequent memory, action, or report, even in the absence of strong new external drive. Therefore current sensation alone neither guarantees nor exhausts conscious content. ■

This proposition is important because it separates the present framework from two common simplifications. The first simplification says that consciousness is just incoming sensation. The second says that consciousness is just abstract integration or complexity. In the present framework, both are incomplete. Consciousness is an operation: internal measurement of a selected and stabilized inner hologram under the Qi-HAL runtime.

One may package the operation in a compact normal form. Let $\mathcal{S}_{\text{DGI}}$, $\mathcal{S}_{\text{GSSB}}$, and $\mathcal{S}_{\text{GA}}$ denote the exploration, selection, and stabilization stages acting on ρ_{IH} . Then the self-accessible content state may be written schematically as

$$\rho_{\text{self}}(t) = \mathcal{M}_{\text{int}} \circ \mathcal{S}_{\text{GA}} \circ \mathcal{S}_{\text{GSSB}} \circ \mathcal{S}_{\text{DGI}}[\rho_{\text{IH}}(t)]. \quad (64.21)$$

Equation (64.21) is the recursive analog of the external measurement chain. It is not a separate metaphysical law; it is the same open-system logic reapplied to the system's internal model.

64.8 Universality across open systems with memory once $U(1)_{Q_i}$ is present

The user asked explicitly that this section show how the framework applies to any open system with memory once $U(1)_{Q_i}$ is applied. We can now state that result clearly.

The key point is that $U(1)_{Q_i}$ makes informational state changes physically budgeted, phase-covariant, and persistent in the same sense as the earlier sections. Once a system has a memory-bearing variable with nonzero informational current J_Q , its state changes can no longer be treated as abstract bookkeeping only. They become lawful gauge-bearing events. This makes an IH possible in a physically disciplined sense rather than as a metaphor.

Because internality is placed only after runtime constitution, the full Qi stack, and persistence/-completeness, the next theorem should be read as a late-stage consequence of the earlier construction rather than as a primitive assumption.

Theorem 64.14 (T-EIC: OH/IH/EIC universality for memory-bearing $U(1)_{Qi}$ systems (principal theorem) (T-Qi4-15.14)). *Let S be an admissible open system satisfying the following:*

- (i) *S has observer-accessible ports and therefore an admissible OH channel;*
- (ii) *S possesses a persistent memory-bearing variable m_t ;*
- (iii) *S carries nonzero informational current J_Q coupled to $U(1)_{Qi}$;*
- (iv) *future observations depend nontrivially on present action or motion, so that a MePMoS loop exists.*

Then:

- (a) *S always admits an outer hologram H_t^{out} relative to the chosen instrument class;*
- (b) *if m_t stores predictive information about future observations, then S admits at least a proto-inner-hologram H_t^{in} ;*
- (c) *if the proto-IH is recurrently re-measured onto an internal boundary and stabilized by $DGI \rightarrow GSSB \rightarrow GA$, then S admits an evolving inner cosmos $\mathfrak{E}_t$;*
- (d) *high-dimensional, multiscale, self-referential instances of $\mathfrak{E}_t$ correspond to what the present framework calls consciousness-grade runtime.*

Proof. Part (a) is immediate from OGGEF and Definition 64.1. Part (b) follows from Corollary 64.6: if a persistent memory variable carries predictive information about future observations, then the system’s internal state is no longer a function only of the current sample and constitutes a lawful latent reconstructive manifold. The presence of $U(1)_{Qi}$ ensures that this state is not merely formal but physically budgeted and phase-covariant, as shown earlier in Parts III and VI. Part (c) follows from Definitions 64.8 and 64.11: once the IH is recurrently measured on an internal boundary and re-enters MePMoS after selection and stabilization, the coupled OH-IH loop is exactly an EIC. Part (d) is the architectural specialization of Part (c) to high-QCI, multiscale, recursively self-modeling systems. ■

The theorem is intentionally graded. It does not force an all-or-none divide between “conscious” and “not conscious.” Instead it yields a continuum:

- (1) **OH only:** systems that can be externally measured but do not store a predictive internal model of the world;
- (2) **OH + proto-IH:** systems with memory that encode some predictive latent structure but do not yet exhibit rich recursive internal measurement;

- (3) **OH + IH + proto-EIC**: systems whose internal model is repeatedly aligned with external evidence and biases future action;
- (4) **consciousness-grade EIC**: systems with sufficiently complete Qi-HAL architecture, multiscale coordination, and self-accessible internal measurement.

This is exactly the kind of graded picture one expects for a universal open-system theory.

Table 43 illustrates the continuum across representative systems.

Table 43: Representative continuum across memory-bearing open systems. Six systems run down the first column — a molecular cycle, a chemotactic bacterium, a ring-attractor navigation circuit, a rodent hippocampo-entorhinal system, a human cortical-hippocampal system, and a distributed institution — with the memory-bearing substrate of each, its likely inner- hologram status, its grade on the coupling, and an interpretation. The table’s structure is its argument: the grades form a continuum rather than a binary, so nothing in the framework requires a threshold at which a system suddenly acquires an inner model.

System	Memory-bearing substrate	Likely IH status	EIC grade	Interpretation
ATP synthase or simple molecular cycle	Conformational occupancy or rotor phase	Degenerate or task-local internal state only	Below proto-EIC	Memory exists, but world reconstruction is minimal and highly local
Chemotactic bacterium	Receptor methylation / adaptive chemical state	Proto-IH of local gradient or resource field	Proto-EIC	Internal memory predicts external chemical structure and guides motion
Ring-attractor navigation circuit	Persistent heading or phase bump	Genuine low-dimensional IH	Intermediate EIC	Internal allocentric heading map coupled to sensory correction and action
Rodent hippocampo-entorhinal system	Distributed recurrent attractor memory	Rich spatial-temporal IH	Strong EIC	Multiscale world model supporting navigation, recall, and imagination
Human cortical-hippocampal system	Large multiscale distributed memory architecture	High-dimensional IH including symbolic models	Consciousness-grade EIC	Internal measurement of the IH becomes reportable, reflective, and abstract
Distributed institution or scientific community	Archived memory, models, communication protocols	Collective IH (distributed world model)	Extended EIC candidate	Group-level reconstruction of external structure through shared memory and measurement

Two examples deserve emphasis.

A chemotactic bacterium. The OH is the local chemical environment sampled through receptors. The memory substrate is the adaptive methylation state. The system does not store a full image of the world, but it stores a predictive estimate of whether the current trajectory is improving or worsening access to nutrient. In this restricted but genuine sense it carries a proto-IH of the local chemical world. The resulting running alignment of sensed gradient and adaptive state is a minimal proto-EIC.

A mammalian navigation system. The OH is the multisensory stream generated by vision, vestibular input, proprioception, and action-dependent re-sampling. The IH is a rich allocentric manifold stabilized by recurrent attractor dynamics, phase coding, and multiscale coordination. The DGI phase corresponds to multiple candidate interpretations or trajectories; the GSSB phase corresponds to selected heading, place, or scene interpretation; the GA phase corresponds to stabilized memory or report. Here the EIC becomes rich enough to support explicit conscious episodes and, in humans, abstract theoretical reasoning.

64.9 Claims, non-claims, and empirical consequences

Because the current section touches the difficult topic of consciousness, it is important to state carefully what is and is not being claimed.

The paper *does* claim that the distinction between external projection and internal reconstruction is constitutionally necessary. It *does* claim that any memory-bearing open system under $U(1)_{Q_i}$ may admit an IH in an appropriately scaled sense. It *does* claim that the EIC is the coupled process aligning inner model and outer evidence. It *does* claim that consciousness should be modeled as internal measurement of a selected and stabilized inner hologram.

The paper *does not* claim that every memory-bearing system is therefore conscious in the same sense as a human. It *does not* claim that the IH recovers the full six-dimensional bulk without quotient loss. It *does not* claim that one single neural frequency or one single anatomical region is “the seat” of consciousness. The framework is architectural and dynamical, not phrenological.

Table 44 summarizes the distinction.

Table 44: Claims and non-claims of the outer/inner hologram section. The two columns are placed side by side because this is the section of the paper where over-reading is most likely and most costly. Everything in the left column is a structural statement about coupled boundary descriptions; everything in the right column is a philosophical claim the framework does not make, including any identification of the construction with subjective experience. The right column is not a disclaimer appended to the left one — it is what makes the left one sayable.

Claim	Non-claim
Every observation is an OH event exported through an admissible boundary channel.	The OH is not the full bulk reality itself.
The IH is a lawful internal reconstruction from partial, often lower-dimensional sensory data.	The IH is not a literal pixel-by-pixel copy of the world.
The EIC is the repeated coupling of OH and IH under MePMoS, Qi-EP, and GFF.	The plus sign in $\text{OH} + \text{IH} = \text{EIC}$ is not ordinary arithmetic addition.
Conscious content is internal measurement of a selected and stabilized IH branch.	Not every internal variable or memory trace is conscious.
Memory-bearing $U(1)_{Q_i}$ systems admit a graded OH/IH/EIC continuum.	The framework does not impose a binary conscious/non-conscious cut at the level of all open systems.

Even at the present drafting stage, the framework also yields empirical consequences. Some are direct, others are programmatic. Table 45 lists the most immediate ones.

Table 45: Representative empirical consequences of the outer/inner hologram framework. Each row is one prediction or consequence, with its content and the place in the section it comes from. The table is deliberately short. Most of the section’s content is organizational, and the discipline this table enforces is that only the rows appearing here are offered as things the world could contradict.

Prediction / consequence	Content	Connection to the section
Closed-loop sensing should outperform passive exposure for latent-structure learning	If motion/action is restricted while sensory drive is preserved, allocentric reconstruction quality should degrade because the observability rank drops.	Theorem 64.5
World-model quality should scale with memory depth and multiscale coordination	Systems with deeper memory and broader BFFT hierarchy should reconstruct richer IH manifolds.	Parts VI, VII, and Theorem 64.14
Internal conscious report should follow selection and stabilization, not raw input onset	Reportable content should lag preconscious activity because DGI exploration must be followed by GSSB and GA.	Definition 64.12
Minimal organisms may show proto-EIC without human-like consciousness	Memory-bearing agents such as bacteria should exhibit predictive internal states aligned with environmental structure, but only in a proto-IH / proto-EIC sense.	Theorem 64.14
Scientific learning is a measurable EIC refinement process	Across training or expertise, internal models should better predict OH data while respecting energetic and entropic cost.	Theorem 64.9

It is useful to compare the present framing with two existing literatures already discussed in the Introduction. Free-energy and active-inference frameworks emphasize prediction, action, and belief updating (Friston, 2010). The present section agrees with that emphasis but locates it inside a larger geometry-to-boundary constitution: the IH update is not a floating agent principle but the internal counterpart of the OGGEp map and the Qi-EP selector. Integrated-information approaches emphasize the structured unity of conscious content (Tononi, 2004). The present section does not reject that intuition; rather, it says that any such unity must be realized as internally measured, selected, stabilized IH structure, not as a primitive postulate independent of runtime, geometry, and ledger.

A final biological point is worth emphasizing. Once one accepts that every observation is OH and that every persistent predictive model is IH, then the developmental and evolutionary history of life can be reread as the progressive enrichment of the EIC. Primitive systems reconstruct only tiny local manifolds. Rich nervous systems reconstruct navigable spatial worlds. Human scientific communities reconstruct abstract lawful structure far beyond immediate sensory appearance. The paper’s hypothesis is that these are not separate miracles but different depths of one common open-system runtime.

64.10 Section summary

The OH/IH/EIC layer is now in place. The main conclusions are:

- (1) The outer hologram is the lawful boundary projection of higher-dimensional or otherwise inaccessible external structure. Every registered observation is an OH event.
- (2) The outer hologram is a quotient object: observation accesses equivalence classes induced by the boundary kernel, not the full bulk state directly.
- (3) The inner hologram is the internal reconstructive manifold built from memory, prediction, action, and sensory history. It is not a stored sensory movie but an allocentric latent model.
- (4) Motion and memory are mathematically necessary for IH formation whenever the instantaneous sensory map is non-injective. This follows from the action-conditioned observability theorem.
- (5) The equation $\text{OH} + \text{IH} = \text{EIC}$ means lawful closed-loop composition, not scalar addition. The EIC is quantified by an alignment functional coupling reconstruction accuracy to the GFF budget.
- (6) In the gradient normal form, the EIC functional is a Lyapunov quantity, so the OH-IH loop monotonically refines internal reconstruction in the noiseless limit.
- (7) Consciousness is not raw sensation and not the whole IH. It is internal measurement of a selected and stabilized branch of the IH on a self-accessible internal boundary under the Qi-HAL runtime.
- (8) Therefore the OH/IH/EIC framework extends GDOS from external measurement to internal world-modeling, conscious content, and the graded continuum of memory-bearing open systems under $U(1)_{\text{Qi}}$.

With this section complete, the manuscript contains the constitutional runtime, the scalar ledger and packet layers, the full Qi-4 pillars, the persistence/completeness vocabulary, and the external/internal holographic semantics needed for the hierarchy, application, universality, discussion, and conclusion chapters that now follow immediately.

Bridge to Part X. With the four pillars and the persistence/completeness/internality framework now established, the question naturally arises: how far does the runtime extend? Part X answers through the five-story pagoda hierarchy theorem and the cross-domain reduction recipe. The same four Qi-4 pillars appear at the floors of particle, cosmos, chemistry, life, and mind, with the differences between domains living only in the SIP data, the selector parameters, and the active sectors — not in the constitutional law. Part X is therefore the existence-and-scope chapter of Qi-4: the runtime is one program, not five.

65 The world one can never reach

Every system in this Part is in the same predicament, and the predicament has no way out. What it can act on is the world; what it can touch is its own boundary. Between the two sits a projection that throws information away, and no amount of instrumentation recovers what the projection discarded, because many different worlds send the same message through.

So the response is not to look harder. It is to build — to maintain something inside that stands in for what cannot be reached, and to keep correcting it against what does arrive. That is what a bacterium is doing when it holds a chemical trace of whether things have been improving. It is what a rodent is doing when it maintains a map that survives its own turning. It is what a laboratory is doing when it holds a theory of six dimensions and checks it against instruments that report four. The difference between these is depth, not kind. All three are the same operation running at different scales, and the last one — science itself — is not an exception to the pattern but its most elaborate instance.

Two consequences follow that are worth carrying out of this Part. The first is that perception is not passive and never was. A system that cannot move cannot, in general, work out what is around it, because the things it might be looking at would all look the same. Moving is how the ambiguity is broken. Knowing and doing are not two faculties, one of which serves the other; they are two phases of a single loop, and the loop does not run in one direction only.

The second is a limit, and it should be stated as plainly as the result. What such a system converges on is never the world. It is an adequate model of a class of worlds — all the ones its boundary cannot tell apart. That is not a defect of biology, or of instruments, or of this framework. It is what the projection does, and it applies to the most careful science exactly as it applies to the bacterium. The honest form of the ambition is not to see things as they are. It is to narrow the class until the difference stops mattering for anything one needs to do.

The Part gives a technical name to the portion of an internal model that has been selected, stabilized and made available to a system's own next cycle. It is a name for something measurable, applied by degree rather than by kind, and it settles nothing about what it is like to be the system

in question. That question is left where it was found. What is claimed is narrower and, if it holds, still worth having: that the same open-system logic which describes how a detector registers a particle, turned inward, describes how a system registers its own state — and that where the loop runs deep enough, this is what a world coming into view from the inside consists of.

Part X

Emergence and scope

Hierarchy, applications, and universality across nature

This is the Part where a framework of this breadth usually stops being science, and the work here is mostly the work of not letting that happen.

The question is how far the runtime reaches. Having built one account of one open system, does the same account describe a molecule, a cell, a brain and a society — and if the claim is that life and mind follow from physics, what exactly is being claimed?

The first move is to refuse to let a level be declared by fiat. Anyone can draw a hierarchy of nature on a napkin; the hierarchies are generally decorative. A level counts as a level here only if it passes three tests that can be applied to a real system: its collective variables nearly close, so that it can be described without constant reference to what is beneath it; it persists far longer than the relaxation of the level below; and it opens at least one genuinely new channel of exchange rather than inheriting all of them from downstairs. Most proposed coarse-grainings fail one of the three, which is why so few levels are scientifically stable. Inside a level, the dynamics is descent on a fixed landscape. Between levels, the ascent is the other regime — a reorganization that opens basins the level below did not possess.

The second move is to say precisely what "inevitable" means, because this is where the framework's boldest sentence lives. Two readings are possible and only one is claimed. It is *not* claimed that every planet, every history, every niche produces life and then mind; that reading is disclaimed in as many words. What is derived is the ensemble statement: in a driven universe whose viable configurations occupy more than a vanishing fraction of the accessible space, a search that keeps returning will reach and hold them with probability one. The determinism is not that every history is the same. It is that the higher levels are attractor classes rather than freak accidents. Two conditions of that argument are named as open.

The Part also settles an old confusion about whether higher levels can act on lower ones. They can, and without any breach of law: what a higher level changes is not the equations underneath but the conditions those equations run under — concentrations, boundary conditions, attention, institutions. Regime modulation, not law violation.


The honest test of the Part is the recipe near its end. It is a seven-step procedure for taking a system from any science and writing it down in this vocabulary, and the framework earns its keep only if a reader in an unfamiliar field can work down that column and get something checkable out of the bottom.

Part [X](#) extends the Qi-4 framework from individual systems to the universal scope it claims. The first section introduces the Japanese five-story pagoda hierarchy — a single ladder running from elementary particles to molecular chemistry, from chemistry to cells, from cells to brains, and from brains to social and technological systems. Section [67](#) walks the cross-domain applications:

catalytic loops, cellular metabolism, neural memory, ecological persistence, and collective intelligence. Section 70 closes the Part with the universality claim — that the same Qi-4 grammar applies to every open system with memory once $U(1)_{Qi}$ is active.

66 Hierarchy and emergence: the Japanese five-story pagoda

In the regime vocabulary of §7, the pagoda is the architecture of the alternation: *within* each story the dynamics is GRIP-In — descent on that story’s fixed landscape, ending at the story’s GGA — while the ascent from one story to the next is a GRIP-Out: a landscape reorganization that opens basins the lower story did not possess. The capstone GG-A11-Origin (Arisaka, 2026, A11-Origin) runs this alternation through eighteen layers; the five-story pagoda below is its load-bearing summary, and the Qi-HAL switch of §52 is the mechanism by which the upper stories manufacture their own GRIP-Out’s.



Pagoda factorization theorem: each higher floor is forced by closure-defect + port-novelty admissibility on the previous.
Conditional inevitability theorem: in a sufficiently rich driven universe, the higher floors are reached almost surely (ensemble sense).

Figure 16: The Japanese five-story pagoda hierarchy of Part X. The five floors — particle, cosmos, chemistry, life, mind — are linked vertically by the pagoda factorization theorem, with the shinbashira central pillar carrying the OGGEp projector and the combined Gi-4/Qi-4 cascade. The choice “five” is the historically natural sequence; whether it is uniquely forced by MSMF on closure-defect and port-novelty admissibility is recorded as open problem O7. The conditional-inevitability theorem (Theorem 66.4) establishes that, in a driven cosmological measure with positive viable subsets, the higher floors are reached almost surely (open-problem O8 records the residual hypothesis bounds).

Part X begins after the theory has reached internality. The next issue is hierarchical emergence: how the same runtime reappears across recurrent floors of natural organization. Section 66 therefore

derives the Japanese five-story pagoda as the hierarchical factorization of the already established runtime.

66.1 Why a universal theory of open systems must become hierarchical

A serious open-system theory cannot stop at one level: a particle, a cosmos, a catalytic network, a cell, and a nervous system are five recurrent levels where the same runtime reappears with different dominant packets, ports, and modeling depth. The Japanese five-story pagoda is therefore introduced as the architectural image of a theorem: persistence in nature is organized as a hierarchy of symmetry-broken floors built around a common invariant spine (DGI) (Maynard Smith and Szathmary, 1995).

Concretely, a new floor at level $\ell + 1$ requires (i) approximate *closure* of collective variables $z^{(\ell+1)} = C_{\ell+1}(z^{(\ell)})$, measured by a small closure defect

$$\Delta_{\ell+1}(t) := \dot{z}_{\text{exact}}^{(\ell+1)}(t) - F_{\text{eff}}^{(\ell+1)}(z^{(\ell+1)}(t)), \quad (66.1)$$

(ii) *persistence*, $\tau_{\text{life}}^{(\ell+1)} \gg \tau_{\text{relax}}^{(\ell)}$, and (iii) *port novelty*, i.e. at least one new effective exchange channel not obtained by pulling forward lower-level ports.

Definition 66.1 (Admissible floor). A proposed higher-level description $\widehat{\text{SIP}}^{(\ell+1)}$ is an *admissible floor* if there exist constants $\varepsilon_{\ell+1} \ll 1$ and $\kappa_{\ell+1} > 1$ such that

$$\langle \|\Delta_{\ell+1}\|^2 \rangle^{1/2} \leq \varepsilon_{\ell+1} \langle \|F_{\text{eff}}^{(\ell+1)}\|^2 \rangle^{1/2}, \quad (66.2)$$

$$\tau_{\text{life}}^{(\ell+1)} \geq \kappa_{\ell+1} \tau_{\text{relax}}^{(\ell)}, \quad (66.3)$$

$$\dim \text{span } \mathcal{P}^{(\ell+1)} > \dim \text{span } \mathcal{P}_{\text{pull}}^{(\ell \rightarrow \ell+1)}. \quad (66.4)$$

Definition 66.1 converts the pagoda metaphor into a checklist: most proposed coarse-grainings fail one of closure, persistence, or port novelty, so only a few scientifically stable floors emerge.

66.2 The pagoda factorization theorem

We now turn the hierarchy claim into a theorem. Consider the total state

$$Z = (z^{(0)}, z^{(1)}, z^{(2)}, z^{(3)}, z^{(4)}) \in \mathcal{Z}^{(0)} \times \mathcal{Z}^{(1)} \times \mathcal{Z}^{(2)} \times \mathcal{Z}^{(3)} \times \mathcal{Z}^{(4)},$$

where the intended interpretation of the five floors is

$$\text{particle} \rightarrow \text{cosmos} \rightarrow \text{chemistry} \rightarrow \text{life} \rightarrow \text{mind}. \quad (66.5)$$

Let the exact reduced flow after bulk-to-boundary descent and gauge reduction be written as

$$\dot{Z} = \mathfrak{F}(Z) + \mathfrak{R}(Z) + \eta(t), \quad (66.6)$$

where $\mathfrak{F}$ is the deterministic drift, $\mathfrak{R}$ is the port-driven reinjection, and η denotes fluctuations.

Assume that for each floor there exists an GFF ledger $\mathcal{F}_{\text{GFF}}^{(\ell)}$, a packet $\widehat{\text{SiP}}^{(\ell)}$, a current bundle $\mathcal{J}^{(\ell)}$, and an admissible selector/attractor pair. Then the flow may be rearranged into the hierarchical block form

$$\dot{z}^{(\ell)} = F_{\text{self}}^{(\ell)}(z^{(\ell)}) + B^{(\ell,-)}(z^{(\ell-1)}) + B^{(\ell,+)}(z^{(\ell+1)}) + \Xi^{(\ell)}(Z) + \eta^{(\ell)}(t), \quad (66.7)$$

with the conventions $B^{(0,-)} \equiv 0$ and $B^{(4,+)} \equiv 0$. The terms have a clear meaning:

- $F_{\text{self}}^{(\ell)}$ is the autonomous floor drift selected by the local packet and local GFF ledger;
- $B^{(\ell,-)}$ is upward inheritance from the immediately supporting floor;
- $B^{(\ell,+)}$ is slow contextual control from the immediately higher floor;
- $\Xi^{(\ell)}$ is the residual non-nearest-floor coupling, which is small when the hierarchy is clean.

This theorem is the first theorem of Part X precisely because hierarchy becomes meaningful only after internality is in place. It translates the already established runtime into a recursive floor structure.

Theorem 66.2 (Pagoda factorization theorem (T-Qi4-16.2)). *Suppose the five floors in Eq. (66.5) are admissible in the sense of Definition 66.1, and suppose further that the residual couplings satisfy*

$$\|\Xi^{(\ell)}(Z)\| \leq \epsilon_{\Xi} \sum_{m=0}^4 \|z^{(m)}\|, \quad 0 < \epsilon_{\Xi} \ll 1. \quad (66.8)$$

Then the full open-system runtime admits a five-floor factorization in which:

- DGI acts as a common gauge spine across all floors;*
- each floor is created by a floor-specific GSSB event selecting an order parameter compatible with the inherited lower-floor packet;*
- each floor is stabilized by a floor-specific GA compatible with the local GFF landscape;*
- upward and downward couplings preserve constitutional consistency in the sense that no higher floor can violate the packet or ledger constraints of the floors below it.*

Proof. We prove the factorization by induction on the floor index ℓ .

Base case ($\ell = 0$). Floor 0 is the constitutional-particle layer: the 6D bulk projected to its 4D boundary via OGGEF. The closure-defect bound of Definition 66.1 is vacuous at floor 0, since there is no floor -1 ; we adopt the convention that floor 0 is admissible by hypothesis, inherited from the constitutional admissibility of the Standard-Model boundary established in Papers I–IV-A. Concretely: an effective packet Π_0 exists (the constitutional state space), an effective ledger F_0 exists (the action ledger of the Standard-Model boundary), a port set $\mathcal{P}_0$ is supplied by the scale-radial dependence, and the master pipeline $\text{DGI} \rightarrow \text{GSSB} \rightarrow \text{GA}$ at floor 0 is Theorem 7.3 of Section 7.

Inductive step. Assume floor ℓ is admissible in the sense above: it has an effective packet Π_{ℓ} , an effective ledger F_{ℓ} , a port set $\mathcal{P}_{\ell}$, and satisfies two distinct floor-creation conditions for the transition to floor $\ell + 1$:

- the *closure-defect bound* $\varepsilon_{\ell+1} \ll 1$, measuring how nearly the would-be floor- $(\ell + 1)$ slow variables form a closed sub-algebra under the floor- ℓ dynamics (this is the smallness parameter);
- the *lifetime separation* $\tau_{\ell+1}/\tau_\ell \geq C_{\ell+1} > 1$, measuring how cleanly the floor- $(\ell + 1)$ timescale separates from the floor- ℓ timescale (this is the separation ratio).

These two conditions play opposite roles: $\varepsilon_{\ell+1}$ controls the size of the residual coupling, and $C_{\ell+1}$ controls the dynamical separation. We construct floor $\ell + 1$ explicitly.

Let $z^{(\ell+1)}$ denote the slow variables: coarse-grained order parameters whose relaxation time exceeds τ_ℓ by the factor $C_{\ell+1}$. Under the closure-defect bound and the lifetime separation, the standard adiabatic-elimination theory of singularly perturbed dynamical systems (Tikhonov’s theorem on slow manifolds, 1952, and its geometric refinement in Fenichel’s invariant-manifold theorem, 1979) supplies an effective generator $\mathcal{L}_{\text{eff}}^{(\ell+1)}$ on the slow manifold that is Lipschitz with constants of order $\varepsilon_{\ell+1}$ (the residual coupling) and that approaches a well-defined limit as $\varepsilon_{\ell+1} \rightarrow 0$ and $C_{\ell+1} \rightarrow \infty$. The strict enlargement of the port space, $\mathcal{P}_{\ell+1} \supsetneq \mathcal{P}_\ell$, ensures that floor $\ell + 1$ couples to environmental channels not already exhausted by lower floors — this is what makes the new floor a legitimate packet rather than a mere coordinate redescription.

DGI at floor $\ell + 1$ (claim (a)). Gauge-equivalent redundancy must be quotiented at every floor. The inherited gauge reduction at floor ℓ propagates upward: gauge-equivalent representations at floor ℓ collapse to identified points at floor $\ell + 1$ by the very definition of the coarse-graining map. Hence DGI acts as a vertical invariant across all floors — the mathematical content of the pagoda’s central spine. A floor that failed to respect the inherited gauge reduction would double-count states already declared equivalent below, violating MSMF and the packet definitions established at floor ℓ .

GSSB at floor $\ell + 1$ (claim (b)). Because the closure-defect bound is satisfied with strict inequality on a non-empty open set of order-parameter configurations, the local $F_{\ell+1}$ landscape admits at least one bifurcation as the floor- $(\ell + 1)$ control parameter is varied. At each bifurcation, GSSB selects a branch. The order parameter selected is the one compatible with the inherited lower-floor packet, because variables outside the floor- $(\ell + 1)$ gauge slice are not available for selection by construction.

GA at floor $\ell + 1$ (claim (c)). The selected branch is stabilized by the local $F_{\ell+1}$ landscape: the lifetime separation $\tau_{\ell+1}/\tau_\ell \geq C$ ensures that the floor- $(\ell + 1)$ dynamics relaxes on a timescale fast compared to its own decay while remaining slow compared to floor ℓ . Standard GA arguments (Theorem 7.2) then apply at the floor- $(\ell + 1)$ level.

Upward and downward consistency (claim (d)). Coupling between floors is mediated by the block-flow matrix $B^{(\ell,+)}$ of Eq. (66.7). Because higher-floor variables are defined as coarse-grained packets of lower-floor variables, a higher floor cannot violate the packet or ledger constraints of any floor below it: any putative violation would correspond to a redefinition of the lower-floor coarse-graining map, in contradiction to the inductive hypothesis. Higher floors can only act on lower floors through slow parameter modulation, boundary conditions, or port scheduling.

Conclusion. Each step of the induction preserves admissibility; the construction terminates at any floor where one of the floor-creation conditions (closure defect, lifetime separation, port-rank enlargement) ceases to be satisfied. In the empirical universe the construction terminates at floor 5 (mind), but the theorem is structural and does not predict the termination point in principle: the count *five* is an empirical reading, not a theorem output (Open Problem O7), the induction itself

yielding an arbitrary-depth tower wherever the floor-creation conditions hold. Like every result built on the imported DGI–GSSB–GA pipeline, the factorization inherits the conditional status of GG-A6-GRIP’s GRIP-cascade closure (conditional on M11 in spectral-flow form; M10 resolved externally, the scope refinements closed at class level). ■

Theorem 66.2 shows that the pagoda is not a vague systems-science metaphor. It is the recursively applied open-system logic of the whole manuscript.

Table 46: The five-story hierarchy as a ladder of admissible floors. The five — particle, cosmos, chemistry, life, mind — run down the first column, with the dominant packet on each floor, the new ports it opens, the outcome the cascade reaches there, and the typical grade of internal model. Read downward, each row’s outcome becomes the next row’s substrate, which is the recursion the framework claims; read across, the columns say what is genuinely new at each level rather than merely larger.

Floor	Dominant packet	Typical new ports	Dominant GSSB/GAoutcome	Typical IH/EIC grade
Particle	fields, quanta, elementary records	scattering, emission, absorption	stable particles, conserved charges, bound elementary structures	OH only or degenerate proto-IH
Cosmos	bound gravitational structures	radiative transport, gravity wells, planetary boundary conditions	stars, galaxies, planets, long-lived gradients	OH plus long-memory environmental scaffolding
Chemistry	atoms, molecules, catalytic loops, compartments	redox exchange, diffusion, catalytic throughput	stable reaction networks, molecular machines, compartments	proto-IH in adaptive molecular systems
Life	cells, organisms, ecosystems with maintained boundaries	metabolism, sensing, motility, repair, reproduction	SeOS persistence, adaptive niches, developmental attractors	strong proto-EIC to intermediate EIC
Mind	nervous systems or equivalent high-QCI architectures	internal measurement, symbolic report, deliberate action, model-sharing	stabilized allocentric IH, self-modeling, conscious episodes	consciousness-grade EIC

66.3 DGI as the pagoda’s central pillar (*shinbashira*)

Traditional Japanese pagodas are famous for their central pillar, the *shinbashira*, which contributes to coherence and resilience without being reducible to any single floor. In the present framework that structural role is played by DGI. The analogy becomes precise if one writes the floor variables as deviations around floor-specific attractors,

$$e^{(\ell)} := z^{(\ell)} - z_*^{(\ell)},$$

and isolates the common invariant variable c carrying the vertical gauge spine. The near-attractor dynamics can be linearized as

$$\dot{c} = f_c(c), \quad (66.9)$$

$$\dot{e}^{(\ell)} = A_\ell(c) e^{(\ell)} + \sum_{m \neq \ell} B_{\ell m}(c) e^{(m)} + \xi^{(\ell)}(t), \quad \ell = 0, \dots, 4. \quad (66.10)$$

If each isolated floor is stable, then A_ℓ has negative spectral abscissa. The question is whether coupling destabilizes the whole tower.

Proposition 66.3 (Shinbashira stability criterion). *Assume that for each floor there exists $\lambda_\ell > 0$ such that*

$$\frac{1}{2}(A_\ell + A_\ell^\top) \preceq -\lambda_\ell I,$$

and define the coupling norm

$$\beta_\ell := \sum_{m \neq \ell} \|B_{\ell m}\|.$$

If $\beta_\ell < \lambda_\ell$ for all floors, then the composite Lyapunov functional

$$V(e) = \sum_{\ell=0}^4 \alpha_\ell \|e^{(\ell)}\|^2, \quad \alpha_\ell > 0, \quad (66.11)$$

can be chosen so that $\dot{V} \leq -\gamma V + C\|\xi\|^2$ for some $\gamma > 0$. Therefore the entire tower remains input-to-state stable.

Proof. Differentiate Eq. (66.11) using Eq. (66.10). The diagonal terms contribute $-2\alpha_\ell \lambda_\ell \|e^{(\ell)}\|^2$. The cross-floor couplings are bounded by Cauchy–Schwarz and Young inequalities:

$$2\alpha_\ell |\langle e^{(\ell)}, B_{\ell m} e^{(m)} \rangle| \leq \alpha_\ell \|B_{\ell m}\| (\|e^{(\ell)}\|^2 + \|e^{(m)}\|^2).$$

Summing over $m \neq \ell$ gives the margin β_ℓ . Because $\beta_\ell < \lambda_\ell$, one can choose positive weights α_ℓ so that the negative diagonal terms dominate the coupling terms. Noise enters additively and yields the bounded forcing term $C\|\xi\|^2$. Hence $\dot{V} \leq -\gamma V + C\|\xi\|^2$. ■

The proposition explains the pagoda analogy in rigorous language. The floors may move relative to one another; they need not be rigidly frozen. What matters is that a common invariant spine plus sufficiently weak cross-floor destabilization preserves the whole tower. In physical and biological language, the invariant spine is the reuse of the same constitutional rules at every level: packet declaration, gauge reduction, current accounting, selector, attractor, and lawful observation.

Table 47: The pagoda analogy translated into runtime language. Each row takes one element of the structure — the central pillar, a floor, the formation of a floor, long-term persistence, the interaction between roofs — and gives its runtime analog and its meaning here. The analogy is retained because it carries one point better than the formal statement does: the floors are not bolted to one another, and what makes the structure stand is that each rests on the one below without being rigidly attached to it.

Pagoda element	Runtime analog	Meaning in the present framework
Shinbashira	DGI	Common gauge-invariant constitutional spine reused across all floors
Each floor	floor-specific packet $\widehat{\text{SIP}}^{(\ell)}$	A persistent effective level with its own ports, observables, and ledger
Floor formation	GSSB	Selection of one organized branch/order parameter among many admissible possibilities
Long-term persistence	GA	Stabilization of the selected branch under the local GFF landscape
Roof and overhang interactions	nearest-floor coupling terms	Slow cross-level boundary modulation without violation of lower-law consistency
Seismic resilience	composite Lyapunov stability	The tower survives perturbation because invariants and margins are distributed rather than localized

66.4 The five floors in detail

We now describe the five floors themselves. The order in Eq. (66.5) is not arbitrary. Each floor supplies the conditions for the next.

Floor 1: particle. This is the floor developed most explicitly in Papers I–IV-A. Its packet contains fields, charges, and the observer-accessible boundary records derived from GG-6 and GDOS. The fundamental current ledger is already locked by the Arizaka Ratio and the Arizaka Equation. The relevant asymptotic structures are stable particles, quasiparticles, and simple bound states. The OH semantics is present at once because measurement is already a projection from inaccessible structure to accessible boundary record. But the internal-model depth at this floor is typically negligible or degenerate: there is usually no rich self-maintained memory architecture.

Floor 2: cosmos. The cosmos floor begins when lower-level matter and radiation organize into long-lived gravitationally structured environments. The simplest instability criterion is the Jeans condition

$$\omega^2 = c_s^2 k^2 - 4\pi G\rho. \quad (66.12)$$

Modes with $k < k_J := \sqrt{4\pi G\rho}/c_s$ are unstable, so perturbations grow. That instability is the cosmic analog of GSSB: one among many homogeneous possibilities is selected. Once cooling, feedback, and orbital regularization operate, stars, galaxies, planetary systems, and persistent gradients appear. Those gradients are essential because higher floors need free-energy throughput

and compartment-friendly boundary conditions. In short, the cosmos floor is the gradient-generating floor.

Floor 3: chemistry. Chemistry begins when stable matter, gradients, and environmental cycling support reaction networks with their own packet identity. A generic reaction-diffusion packet obeys

$$\dot{n}_i = \sum_r \nu_{ir} R_r(n, m) - \nabla \cdot J_i + \xi_i, \quad (66.13)$$

possibly augmented by a conformational or catalytic memory variable m . The chemistry floor is not just “matter at lower temperature.” It is the floor at which event identities, catalytic loops, compartments, and stoichiometric budgets become the natural description. Here Qi-QUA and GFF become especially transparent: reaction events carry energy and information cost simultaneously, and the chemical natural-unit Arisaka Equation becomes operationally useful.

Floor 4: life. Life begins when a chemical packet becomes a SeOS. The defining step is not merely self-replication or metabolism in isolation but packet-level maintenance under continued openness. Using the notation of Sections 8 and 9, and of Part VII, one requires a positive average maintenance margin,

$$\langle \Pi_{\text{maint}} - \mathcal{R}_{\text{repair}}^{\min} \rangle > 0, \quad \langle \dot{\mathcal{F}}_{\text{GFF}} \rangle_{\text{cycle}} \leq 0, \quad (66.14)$$

with boundary integrity and memory depth sufficient to guide action or regulation. The life floor therefore adds adaptive sensing, active repair, motility or transport control, and eventually reproduction and evolution. Proto-IH and proto-EIC are common here: the system carries an internal state that predicts environmental relevance strongly enough to bias future behavior.

Floor 5: mind. Mind begins when the internal model is not merely useful for immediate control but becomes a high-dimensional, recurrently re-measured, self-accessible inner hologram. In the language of Section 64, the system must admit a nontrivial internal boundary and a stabilized selected content. A minimal criterion is

$$\dim H_t^{\text{in}} \gg 1, \quad \text{QCI}(t) > \text{QCI}_{\text{crit}}, \quad \text{rank } \mathcal{B}_{\text{BFFT}} \geq 2, \quad (66.15)$$

where QCI is the Qi-4 Completeness Index of Part VII and $\mathcal{B}_{\text{BFFT}}$ denotes the multiscale tree of frequencies or timescales supporting coordinated memory. This floor supports explicit navigation, planning, imagination, and in human cases symbolic abstraction and scientific theorizing.

The five floors are summarized in Table 48.

Table 48: Operational criteria distinguishing the five floors. Each row is one floor, with the new variable it introduces, the minimum criterion a system must meet to be on it, its memory substrate and its typical stabilizer. The third column is what makes the hierarchy checkable rather than descriptive: it states, for each floor, a condition that can in principle be evaluated for a given system, so that floor membership is a question with an answer.

Floor	New variable	Minimum criterion	Memory substrate	Typical stabilizer
Particle	charges, fields	gauge-consistent packet and current ledger	none or local phase	conservation laws, bound states
Cosmos	gravitational gradients, orbits	persistent large-scale gradients	orbital, environmental	gravity, radiative feedback
Chemistry	stoichiometric loops, catalysis	reaction-network closure with compartments	conformational, catalytic	kinetic attractors, compartments
Life	maintained boundary, repair budget	positive maintenance margin (SeOS)	genetic, epigenetic, regulatory	adaptive regulation, re-production
Mind	self-accessible allocentric IH	high-dimensional IH, internal measurement	neural recurrent memory	recurrent attractors, reportable EIC

66.5 Probabilistic recurrence of high-floor packets in a driven open universe (conditional)

The user-level slogan of the present framework is that life and mind emerge naturally, deterministically, and inevitably in a six-dimensional open universe. To turn that slogan into a mathematical claim, one must distinguish two meanings of inevitability.

The first is *trajectory inevitability*: every local history, on every planet or in every niche, produces the same outcome. The present paper does *not* claim that. The second is *ensemble inevitability*: once the lower floors generate sufficiently many driven windows and the viable higher-floor packets occupy a set of positive measure, the open-system search process reaches and stabilizes those packets with probability one in the long-time, many-trial limit. It is this second meaning that the theory derives.

Let $W_{\text{life}} \subset \mathcal{Z}^{(3)}$ be the subset of chemistry-floor packets that satisfy the SeOS inequalities for sustained maintenance, and let $W_{\text{mind}} \subset \mathcal{Z}^{(4)}$ be the subset of life-floor packets that satisfy the mind criterion (66.15). Let τ_W denote the first hitting time of a set W . Assume that the DGI exploration process is recurrent or sufficiently mixing on the accessible driven region and that every entry into a viable window has a positive stabilization probability under GSSB plus GA.

Theorem 66.4 (Conditional inevitability theorem (T-Qi4-16.4)). *Assume:*

- (i) *the driven accessible region of state space has finite invariant measure μ ;*
- (ii) *$\mu(W_{\text{life}}) > 0$ and $\mu(W_{\text{mind}}) > 0$;*
- (iii) *the exploration process returns infinitely often to every positive-measure open set in the accessible region;*

(iv) conditional on entering W_{life} or W_{mind} , the probability of stabilization by GSSB plus GA is bounded below by positive constants $p_{\text{life}}, p_{\text{mind}} > 0$.

Then

$$\mathbb{P}(\tau_{W_{\text{life}}} < \infty) = 1, \quad \mathbb{P}(\tau_{W_{\text{mind}}} < \infty \mid \tau_{W_{\text{life}}} < \infty) = 1. \quad (66.16)$$

In this precise sense, life and mind are inevitable attractor classes of a sufficiently rich driven universe.

Proof. By assumption (iii), the process enters each positive-measure viable set infinitely often almost surely. Consider first W_{life} . Let E_n be the event that the n -th visit to W_{life} fails to stabilize. By assumption (iv), $\mathbb{P}(E_n \mid \text{visit}_n) \leq 1 - p_{\text{life}}$. Therefore the probability that stabilization fails on the first N visits is bounded by

$$\mathbb{P}(E_1 \cap \cdots \cap E_N) \leq (1 - p_{\text{life}})^N.$$

Taking $N \rightarrow \infty$ yields zero. Hence stabilization in W_{life} occurs almost surely. The same argument applies to W_{mind} conditioned on the existence of life-floor packets from which the mind-floor accessible region is explored. This proves Eq. (66.16). ■

The theorem clarifies the meaning of “deterministic” in the present program. The determinism is not that every micro-history is identical. The determinism is that the runtime plus the measure-theoretic structure of the accessible universe makes higher-floor packets attractor classes rather than freak accidents. Once this is understood, the otherwise bold claim that life and mind are natural and inevitable becomes mathematically disciplined.

66.6 Upward inheritance and downward modulation

A common confusion in multiscale science is that a higher level must either be reduced away entirely or else be granted miraculous causal powers. The present framework avoids both extremes. Upward emergence is encoded by coarse-graining and selection of persistent packets. Downward influence is encoded only through lawful modulation of lower-level boundary conditions, control parameters, and port schedules. A higher floor never violates a lower floor’s law; it changes the regime in which the lower floor operates.

Mathematically one may write the lowest-floor dynamics as

$$\dot{z}^{(0)} = F^{(0)}(z^{(0)}; \theta^{(1)}(z^{(1)}), \theta^{(2)}(z^{(2)}), \dots, \theta^{(4)}(z^{(4)})) + \eta^{(0)}. \quad (66.17)$$

The functions $\theta^{(\ell)}$ are slow control parameters: concentrations, membrane states, tissue architecture, attention, or institutions depending on the floor. Equation (66.17) is the correct formal meaning of downward causation in the present framework. It is regime modulation, not law violation.

66.7 Section summary

The hierarchy chapter establishes five points.

- (1) A higher scientific floor is legitimate only if it closes approximately, persists long enough, and introduces genuinely new effective ports.

- (2) Under those conditions the universal runtime recursively factorizes into the five-story pagoda: particle, cosmos, chemistry, life, and mind.
- (3) DGI plays the role of the pagoda's *shinbashira*: a common invariant spine reused across all floors.
- (4) GSSB forms the floors by selecting specific organized branches, while GA stabilizes those branches under the relevant GFF ledger.
- (5) If driven viable windows for life and mind have positive measure and are repeatedly explored, then life and mind become inevitable attractor classes in the ensemble sense proved by Theorem 66.4.

The hierarchy is now in place. We can therefore move from formal runtime to explicit cross-domain application without losing the floor structure that tells us how those applications are organized.

67 Cross-domain applications: chemistry, life, brains, and society

Once the hierarchy is explicit, the manuscript can move from architecture to worked reduction. Section 67 shows how the same packet-ledger-selector-diagnostic-architecture-internality vocabulary is instantiated in chemistry, life, brains, and society. The order matters: applications now come after the formal hierarchy rather than being used to motivate it prematurely.

67.1 A canonical reduction recipe for any open system with memory

The previous sections supplied the universal language. An application section must now answer a practical question: given a concrete system in any branch of science, how does one actually reduce it into the GDOS–Qi runtime? The answer is a seven-step map.

Definition 67.1 (Universal application tuple). For an open system S with memory under nontrivial $U(1)_{Qi}$, define its application tuple by

$$\mathfrak{R}(S) := \left(\widehat{\text{SIP}}_S, \mathcal{F}_{\text{GFF}}^{(S)}, \mathcal{J}_S, \mathfrak{H}_S^{\text{out}}, \mathfrak{H}_S^{\text{in}}, \mathfrak{E}_S, \Pi_S \right), \quad (67.1)$$

where $\widehat{\text{SIP}}_S$ is the canonical packet, $\mathcal{F}_{\text{GFF}}^{(S)}$ the scalar ledger, $\mathcal{J}_S$ the locked currents, $\mathfrak{H}_S^{\text{out}}$ the OH object, $\mathfrak{H}_S^{\text{in}}$ the IH object, $\mathfrak{E}_S$ the EIC object, and Π_S the floor placement in the pagoda hierarchy.

The tuple is useful only if it always exists under the hypotheses of the paper. After the hierarchy theorem, the next theorem turns the architecture into a reusable reduction method for concrete domains.

Theorem 67.2 (Cross-domain reduction theorem (T-Qi4-17.2)). *Let S be an open system with:*

- (i) *explicit ports and non-isolated throughput,*

- (ii) a persistent memory-bearing state variable,
- (iii) a nontrivial informational current coupled to $U(1)_{Q_i}$, and
- (iv) a chosen instrument class defining external observables.

Then the application tuple $\mathfrak{R}(S)$ in Eq. (67.1) exists up to gauge-equivalent packet reparameterization.

Proof. A packet exists by Section 9: every admissible open-system analysis begins by declaring a SIP. The scalar ledger and current bundle exist by Section 8 and Part III. The selector and architecture exist by the Qi-EP and Qi-HAL Parts, respectively. Because an instrument class is given, the OH object exists by OGGEF and Section 64. Because the system is memory-bearing and carries nontrivial $U(1)_{Q_i}$, the universality theorem of Section 64 implies the existence of at least a proto-IH; if internal measurement is sufficiently structured, an EIC object exists. Finally, the hierarchy chapter supplies the floor placement. Gauge-equivalent packet choices produce equivalent reduced descriptions because DGI identifies redundancy. ■

Table 49 rewrites the theorem as a practical workflow.

Table 49: Canonical reduction recipe for applying the framework to any open system with memory. Each row is one step, with the question to ask and the output it produces. The recipe is included because a framework of this breadth is only useful if it can be run on a case, and the honest test of the paper is whether a reader in an unfamiliar domain can work down this column and produce something checkable. The steps are ordered so that each output is the next step’s input.

Step	Question	Output
1	What is the admissible system packet?	$\widehat{\text{SIP}}_S$: state manifold, ports, observables, ledger bundle
2	What currents flow through the packet?	$\mathcal{J}_S$: locked event currents and dominant throughputs
3	What scalar budget drives lawful relaxation and maintenance?	$\mathcal{F}_{\text{GFF}}^{(S)}$ and maintenance balance laws
4	What trajectories are admissible or selected?	Qi-EP sector: PLA/PLT/MEPP/FEP interpretation
5	What scale signatures diagnose the runtime?	Qi-CON sector: power laws, fractality, two-knee spectra, scale covariance
6	What architecture stores memory and forms world-models?	Qi-HAL sector: MePMoS, NHT, HAL, BFFT, OH/IH/EIC
7	Where does the system sit in the pagoda?	Floor placement Π_S and cross-floor couplings

The value of the recipe is that it does not change across sciences. Only the variables entering each slot change.

68 What each floor is answerable for

A hierarchy that runs from particles to minds is the paper’s largest claim about its own reach, and a claim of that size earns the question of what, exactly, it commits to at each level. This section

answers it floor by floor: which pillar result is load-bearing there, at what grade it is carried, and what observation would refute the reading at that floor specifically. The purpose is to prevent the pagoda from being read as a metaphor with five levels — the failure mode of every hierarchy of this shape in the literature — by making each floor separately answerable.

Three cautions govern the table. First, the grades are not uniform and are not meant to be: the lower floors inherit theorem-grade results, and the upper floors inherit those same results plus placements that the paper marks as organizational rather than predictive. Second, a floor’s entry names what would refute *the reading at that floor*, not what would refute the framework: the pillar falsifiers sit upstream and are stated in Part XI. Third, the floor count itself is an empirical reading and not a theorem output, which the open-problem list records; nothing in the table below turns on there being exactly five.

Table 50: The per-floor accounting of the hierarchy. Each row names the pillar content that is load-bearing at that floor, the grade at which the paper carries it, and the observation that would refute the reading *at that floor*. The grades fall as the floors rise, which is the honest shape of a construction that is theorem-grade where it prices events and organizational where it places behaviors: a reader who wants only what is proved can stop after the second row, and the table says so rather than leaving the boundary to be inferred.

Floor	What is load-bearing there	Grade	What would refute the reading there
Particle	The four-current lock and the single conversion constant; elementary events as the unit of the ledger	Derived	An elementary event of the channel whose four-way bookkeeping departs from the single conversion factor beyond the stated tolerance
Cosmos	The same lock at cosmological port scales; the runtime’s three moves on the geometric companion’s moduli	Derived, imported	A cosmological boundary budget requiring a second conversion factor, or a descent inconsistent with the imported cascade
Chemistry	The current gauge and its two readings; reaction networks as sustained throughputs under clamped ports	Derived, with the failure boundary of Theorem 30.1	A driven reaction network selected by an entropy-production extremum where the circulating sector demonstrably does ledger work
Life	The memory architecture and the persistence layer; the informational column priced at Q_B , anchored by the bond that carries two quanta	Derived; the anchor measured to 1.6%	A metabolic anchor disagreeing with $2Q_B$ beyond calorimetric precision with no integer-multiplicity reinterpretation available
Mind	The internal-measurement loop and the inferential gauge, the latter conditional on an explicit rank condition	Conditional; the floor count itself an empirical reading	Inference maintained where the rank condition provably fails, or an identified inference metric that is not the one the ledger fixes

Read down the grade column and the paper’s shape appears: theorem-grade where it prices events and derives selection, conditional where it requires a checkable structural condition, and organizational where it places behaviors that no equation here predicts. That descent is not a weakness to be hidden in a late section; it is the reason the hierarchy is worth stating at all. A

construction that claimed the same grade at every floor would be claiming that mind follows from the ledger as tightly as the conversion constant does, and it does not.

69 The exportable result: inference derived, not assumed

One result of this paper travels further than the rest, and a reader who came for the biology may reasonably never open the pillar Part in which it is proved. It is worth stating on its own, in the register of the sciences that will use it.

In the study of brains and adaptive systems there is a widely used principle holding that anything which persists behaves so as to minimize a quantity measuring the mismatch between what it expected and what arrived. As mathematics it is productive; as a claimed law about matter it has had a difficult history, because the conditions under which a lump of matter can be said to carry a model of its surroundings turn out to be special rather than generic, and the principle's own defenders now describe it as a way of modeling systems rather than a fact about them.

What Part IV establishes is that the working core of that principle is not a postulate here at all. It is what one obtains by taking the boundary functional this program derives from the six-dimensional action — a functional that knows nothing of brains, and was constructed to answer a question about transport — and asking it which *beliefs* a system updates, rather than which path it takes or which currents it carries. The answer that falls out is natural-gradient descent of a free energy, with the geometry of the descent fixed by the ledger's own informational term rather than chosen for convenience. Three consequences follow, and each is checkable.

The metric stops being a modeling choice. In the usual presentation the geometry on the space of beliefs is installed by hand and the weighting of precisions is a family of free parameters. Here the metric is the second variation of the ledger's informational column, which means that where it can be measured, it is a prediction rather than a fit.

The boundary of the system stops being a hopeful assumption. Nothing here asks that a statistical boundary be discoverable inside the stationary density of a generic system. The split between a system and its surroundings enters as declared data — the ports of the packet, written down with the system and carrying the conditions under which the description would be judged wrong.

The existence of an internal model becomes a condition one can test. Whether a given system admits the inference reading is decided by whether an explicit action-conditioned observability matrix has full column rank. Systems that fail it are classified rather than assumed away, which is the difference between a framework that explains the negative cases and one that is embarrassed by them.

And one caution belongs in the same breath, because a result that arrives with its limits stated is worth more than one that does not: the inference reading fails exactly where the entropy readings fail, in the circulating sector of Theorem 30.1, and the prospective policy machinery built on top of the principle in its own literature does not follow from anything derived here and is not claimed. What is offered is narrower than the original principle and, for that reason, more useful: a derivation of its working core, with the conditions printed on it.

70 Why GDOS applies to all open systems in nature

The universality claim now follows in one paragraph. Once a system is open it has ports; once it persists it requires a packet and a maintenance ledger; once it stores predictive information it requires a memory architecture and a nontrivial informational current; once that current is physically budgeted by $U(1)_{Qi}$ the system pays a lawful cost for carrying and updating memory. From that moment the entire runtime—DGI/GSSB/GA, GFF, SIP, Qi-QUA/Qi-EP/Qi-CON/Qi-HAL, SeOS/QCI, and OH/IH/EIC—is available. None of these ingredients depends on whether the substrate is molecular, biological, neural, or social: only the state variables, timescales, and observables change. This completes the derivational arc, so Part XI can now separate assessment into significance/comparison and limitation/falsifiability.

Bridge to Part XI. Having established the four pillars of Qi-4, the persistence-and-internality framework of Part IX, and the pagoda hierarchy of Part X, what remains is the broader assessment: how does Qi-4 compare to existing programs in open-system science; what are the limitations of the present derivational stage; what are the named experimental falsifiers; and what does the framework imply, methodologically and philosophically, about the open-system world? Part XI gathers the discussion, the honest L6 open-problem triage, the empirical falsifiability program, and a closing meditation on the four pillars as one geometric object.

71 One grammar, many sciences

The pagoda in this Part is a real building, and the reason it is used here is not the shape but the engineering. Those towers survive earthquakes that flatten more rigid structures, and they do it because the floors are not bolted to one another. Each rests on the one below. A central pillar runs the height of the building and belongs to no floor at all. Under load the stories move against each other, and the movement is what saves them.

That is the claim about nature this Part is making. Chemistry is not bolted to particle physics, and biology is not bolted to chemistry. Each rests on what is beneath it, is constrained by it, and cannot violate it — and each is free in the ways that matter, because what a higher level changes is never the law below but the conditions under which that law is running. Concentrations, boundary conditions, attention, institutions: these do not overrule anything. They set the regime. The old dilemma that a higher level must either be reduced away entirely or granted powers that break physics turns out to be a false one, and the resolution is unglamorous. It was never a contest of causes; it was a difference of timescale.

The word "inevitable" has been used about life in this program, and this Part is where it is finally pinned down, so it should be said once more in the plainest terms. It does not mean that any given world will produce living things. Nothing here says that of the Earth, and nothing here would be embarrassed by a sterile galaxy. What it means is a statement about a search rather than about a history: that in a universe which keeps generating gradients, and in which the configurations capable of holding themselves together occupy more than a vanishing fraction of what is reachable, such configurations are not accidents waiting on improbable luck. They are where the search ends

up. And the number of stories in the tower is not a prediction of this theory — five is what has been observed, not what has been derived, and the theory is honest enough to say so.

What would it mean if the grammar really is one grammar? Not that a society is a large molecule, which it is not, or that a cell can be computed from quarks, which it cannot. Something more modest and more useful: that the questions worth asking of an open system are the same questions at every scale. What does it spend. What does it hold on to. What would kill it. Where does it stop. Biology, chemistry and physics would remain separate sciences, because their variables and timescales and instruments are genuinely different — but they would stop being separate *kinds* of explanation. The differences would be differences of subject, which is what they always looked like from a distance, and never differences of law.

Part XI

Assessment

Significance, limitations, falsifiability, and conclusions

72 Discussion: significance, uniqueness, and comparison with existing frameworks

Part [XI](#) begins only after derivation, hierarchy, and applications are complete. The first assessment question is positive and comparative: what is genuinely new, why does it matter, and how does it relate to existing scientific programs? Section [72](#) answers that question before any limitation or falsifiability claims are stated.

72.1 What is new in the present paper

The present manuscript does not claim novelty merely by combining familiar words such as geometry, entropy, information, life, and consciousness. Its stronger claim is structural. It gives one continuous derivational chain from constitutional axioms to a practical runtime for open systems. That chain was not available before the four papers are read together.

The main novelties of the present paper are the following.

- (1) Open systems are taken as fundamental rather than as corrections to ideal closed systems.
- (2) Observation is treated constitutionally as bulk-to-boundary projection, not as a late interpretive patch.
- (3) Information is given a gauge-current status through $U(1)_{Q_i}$ and the Arisaka ledger rather than remaining metaphorical.
- (4) The same runtime vocabulary is reused from particles to society through packet, ledger, selector, diagnostic, architecture, persistence, and holographic layers.
- (5) Consciousness is placed late in the derivation as internal measurement of selected **IH** content rather than early as an unexplained primitive.

These moves are jointly stronger than any one of them separately.

Table 51: Why the framework is structurally distinct from standard one-layer approaches. Six neighboring programs run down the first column, from open quantum systems to hierarchy and complexity theory. The second column says what each captures well — and is meant to be read as credit, not as setup — and the third what the present framework adds. The distinction claimed throughout is not that the neighbors are wrong but that each works at one layer, where the object here is the relation between layers.

Approach	What it captures well	What the present framework adds
Open quantum systems	reduced dynamics, decoherence, instruments	embeds those tools into a larger geometric and multiscale runtime extending beyond quantum experiments (Breuer and Petruccione, 2002)
Nonequilibrium thermodynamics	dissipation, transport, entropy production	adds explicit packet identity, current locking, memory architecture, and lawful observation (Nicolis and Prigogine, 1977, Prigogine, 1967)
Free-energy / active inference	prediction, action, belief update	derives such selector behavior inside a larger packet-ledger-gauge constitution (Friston, 2010)
Integrated-information style theories	structured conscious integration	relocates consciousness after packet, ledger, selector, architecture, and measurement semantics (Tononi, 2004)
Synergetics / self-organization	order parameters and pattern formation	supplies a complete bridge from fundamental geometry to multiscale attractors (Haken, 1983)
Hierarchy/complexity viewpoints	architecture of levels and emergence	provides explicit closure, persistence, and port criteria for legitimate floors (Anderson, 1972, Simon, 1962)

72.2 Relation to existing scientific programs

It is helpful to state the relation to several existing programs one by one.

To nonequilibrium thermodynamics and dissipative structure theory. The present framework agrees that far-from-equilibrium flux can generate organized structure (Nicolis and Prigogine, 1977, Prigogine, 1967). It adds a universal packet, a current-locking ledger, and an explicit hierarchy of observation and internal modeling.

To active inference and predictive processing. The present framework agrees that perception and action should be treated jointly (Friston, 2010). It adds an earlier constitutional layer: prediction is not the first principle but a boundary manifestation of the deeper geometry-to-ledger runtime.

To synergetics and complexity science. The present framework agrees that order parameters, multiscale hierarchy, and emergence matter (Anderson, 1972, Haken, 1983, Simon, 1962). It adds explicit closure-defect, lifetime, and port criteria, allowing one to say when a level is genuinely admissible rather than merely evocative.

To consciousness science. The present framework agrees that conscious content involves structured integration and internal availability (Tononi, 2004). It adds a derivational path from packet to internal measurement, placing consciousness after rather than before the runtime.

73 Cross-Paper Coherence Map

The Qi-4 framework imports its constitutional inputs from Arisaka (2026, A1-Constitution), bulk geometry from Arisaka (2026, A2-GG6), anomaly closure from Arisaka (2026, A3-BI6), boundary-projection structure from Arisaka (2026, A5-GDOS), and its geometric pillar template from Arisaka (2026, A6-GRIP); in turn it exports the four Qi-pillars, SeOS/EIC framework, and Arisaka constant to Arisaka (2026, A9-SM), Arisaka (2026, A10-Cosmos), Arisaka (2026, P1-KK), and Arisaka (2026, C1-Kaxion). Each row of the maps below names a specific theorem number, section, or appendix location in the cited companion paper¹.

73.1 Inward map: imports into GG-A7-Qi4 from the foundation block

The following longtable lists, for each theorem-level fact that Qi-4 imports from a foundation-block paper, the specific companion-paper location and the role the imported fact plays in the present manuscript.

Table 52: Inward coherence map: what this paper imports from the foundation block. Each row is one imported fact, with the companion paper and location that establishes it and the role it plays here. Together with Table 53 it is the paper’s full interface with the cohort, and it is the table to consult when checking that no argument here is circular: everything in the middle column is upstream, and nothing in it depends on a result of this paper.

Imported fact	Companion-paper location	Role in GG-A7-Qi4
Four Axioms AX-1–AX-4 (LC, GGEP, OGGEF, MSMF) and seven postulates	GG-A1-Constitution (Arisaka, 2026, A1-Constitution)	Constitutional source of the four Qi-pillars; AX-3 OGGEF is the foundational dynamics from which the boundary projection underlying Qi-4 is derived.
Arisaka equation $J_E/Q_B = J_e = J_B = J_{Q_i}$	GG-A1-Constitution (Arisaka, 2026, A1-Constitution)	Inherited as the central boundary identity that Qi-QUA (Pillar 1) formalizes through the rank-one inflow theorem and the natural-unit equivalence.
6D bulk metric and gauge group $\text{Spin}(1, 5) \times \text{SU}(5) \times \text{U}(1)_B \times \text{U}(1)_{Q_i}$	GG-A2-GG6 (Arisaka, 2026, A2-GG6)	Bulk arena and four-factor gauge structure from which the four boundary currents (J_e, J_B, J_{Q_i}, J_E) of Qi-QUA descend.

continued on next page

¹Audit trail: GG-A1-Constitution (Axioms AX-1–AX-4; Postulates POS-1–POS-7; Arisaka equation $J_E/Q_B = J_e = J_B = J_{Q_i}$). GG-A2-GG6 (canonical 6D metric; BI-6 four-component vector; OGN-5 cascade (1, 3, 5; 7, 35)). GG-A3-BI6 (Theorem 3.4 OGN-5 trace-ratio lock; Theorem 4.1 anomaly freedom). GG-A5-GDOS ($\mathcal{D}_{\text{OGGEF}} = \mathcal{R} \circ \text{Tr}_{\text{env}}$; SIP; OH+IH=EIC). GG-A6-GRIP (principal theorems T1–T4 with compatibility propositions $\mathcal{C}_{\text{TGG}}, \mathcal{C}_{\text{CCI}}, \mathcal{C}_{\text{GFF}}, \mathcal{C}_{\text{GRIP}}$). Each row names a specific theorem identifier rather than a description, so a change in a companion shows up as a mismatch rather than as a silently wrong pointer.

(continued from previous page)

Imported fact	Companion-paper location	Role in GG-A7-Qi4
OGN-5 Diophantine cascade ($M, N, L; J, K$) = (1, 3, 5; 7, 35)	GG-A2-GG6, Theorem 6 (Arisaka, 2026, A2-GG6)	Sets the integer $K = 35$ used in the Arisaka constant $Q_B = (2/5) E_{\text{Ryd}}/K$; the open problem O5 (OGN-5 selection rule) is precisely the question of whether this assignment is uniquely forced by MSMF.
BI-6 closure $dH_3 = X_4$ with coefficient vector (1, -1, -2/5, -2/5)	GG-A3-BI6, Theorem 3.4 (Arisaka, 2026, A3-BI6)	Provides the rank-one bulk source X_4 from which the four-current lock of Qi-QUA descends. The defensive uniqueness clause cites § 8.7 here.
GS-as-GGEP “geometrization reading”	GG-A3-BI6, § 8.7 (Arisaka, 2026, A3-BI6)	The recap-bridge section of Part I cites this for supportive framing of the Green–Schwarz construction; the present paper does not re-derive the argument.
Anomaly-completion and unified gauge consistency	Arisaka (2026, A3-BI6) (formerly GD-2026-5)	Cross-reference for the legality conditions imported into both GG-A6-GRIP and the present Qi-runtime paper.
GDOS boundary law $\mathcal{D}_{\text{OGGEP}} = \mathcal{R} \circ \text{Tr}_{\text{env}}$ as CPTP	GG-A5-GDOS (Arisaka, 2026, A5-GDOS)	Boundary-projection construction that underlies the operational layer of all four Qi-pillars (Qi-QUA, Qi-EP, Qi-CON, Qi-HAL); the OH+IH=EIC framework in Part IX descends directly from this.
SIP (System Instance Packet) construction	GG-A5-GDOS (Arisaka, 2026, A5-GDOS)	Provides the bookkeeping primitive for the persistence class SeOS (Part IX) and the QCI completeness index.
TGG (Total Geometric-Gauge object) + Master Conservation Law	GG-A6-GRIP, Theorem (Arisaka, 2026, A6-GRIP)	T1 Geometric backbone of Qi-QUA: the rank-one inflow theorem is the boundary face of T1’s master conservation law.
CCI (Covariant Conservation Identity) + OGN-5 uniqueness	GG-A6-GRIP, Theorem (Arisaka, 2026, A6-GRIP)	T2 Geometric scaffolding for the four-current Ward package used in Qi-QUA and Qi-EP.
GFF (Geometric Free-Energy Functional) with $F_{\text{GFF}} = F_{\text{GEP}} + F_{\text{CEP}}$	GG-A6-GRIP, Theorem (Arisaka, 2026, A6-GRIP)	T3 Inherited as the master functional from which Qi-EP extracts PLT, MEPP, and FEP as three boundary gauges (common-origin theorem of).
GRIP cascade $\text{DGI} \rightarrow \text{GSSB} \rightarrow \text{GA}$	GG-A6-GRIP, Theorem (Arisaka, 2026, A6-GRIP)	T4 Inherited as the runtime structure on the family-projector + memory moduli that Qi-HAL formalizes as MePMoS / NHT / HAL / BFFT.
SeOS (Selfish/Self-maintaining Open System) seed	GG-A5-GDOS (Arisaka, 2026, A5-GDOS)	The persistence-class seed that Part IX of the present paper formalizes into the persistence criterion + QCI-completeness index.

73.2 Outward map: exports from GG-A7-Qi4 to downstream papers

The following longtable lists, for each theorem-level fact that Qi-4 exports to a downstream paper, the present-paper theorem label or section location and the downstream paper’s anticipated use.

Table 53: Outward coherence map: what this paper exports to the papers downstream of it. Each row is one exported result, with its location here and the downstream paper that uses it. Read beside Table 52 it shows the paper’s position in the cohort’s dependency graph — a consumer of the foundation block and a supplier to the phenomenology — and it is what a reader auditing the cohort for circularity should read first.

Exported GG-A7-Qi4 fact	A7 location	Downstream paper and anticipated use
Qi-QUA rank-one inflow theorem	Theorem 12.2 (T-QUA.1)	GG-A8-QM: rank-one boundary source structure as the geometric content of quantum measurement events. GG-A9-SM: rank-one ledger for the 28 Standard-Model parameter derivations.
Arisaka constant $Q_B \approx 0.1555 \text{ eV/qubit}$	Theorem 14.3, Eq. (14.5)	GG-A9-SM: energy–information conversion factor for the lepton/quark mass ratios. GG-P1-KK: high-energy scaling factor for KK resonance lifetime predictions. GG-C1-Kaxion: low-energy conversion factor for Kaxion DM detection thresholds.
Compatibility proposition $\mathfrak{C}_{QUA}$	Part III, eq. (10.1)	GG-A8-QM: dual-layer template for the algebraic-vs-operational structure of quantum measurement events.
Qi-EP common-origin theorem	Theorem 32.1 (T-EP.1)	GG-A9-SM: variational backbone for the seesaw-mechanism derivation and the cosmological-evolution selectors of GG-A10-Cosmos.
Three named gauges PLT, MEPP, FEP	Theorems 27.1, 28.1, 29.1	GG-A8-QM: three falsifiable laboratory predictions on action minimization, current statistics, and inferential model quality. GG-A10-Cosmos: cosmological-entropy production as the MEPP gauge on the boundary of the 6D bulk.
Qi-CON operational conformality theorem	Theorem 39.1	GG-A9-SM: conformal structure on the lepton/quark mass moduli. GG-P1-KK: spectral-shape predictions for KK resonance clusters.
Qi-CON two-knee spectral theorem	Theorem 42.1 (T-CON.2)	GG-A8-QM: $1/f$ noise predictions for quantum-measurement decoherence. GG-C1-Kaxion: low-mass DM signal-shape prediction. (Open problems O5–O7 record the residual gaps.)
Qi-HAL universality theorem	Theorem 53.1 (T-HAL.1)	GG-A8-QM: memory architecture underlying any quantum-measurement device. (“Unique” is weakened to “sufficient” here; uniqueness is open problem O6.)

continued on next page

(continued from previous page)

Exported GG-A7-Qi4 fact	A7 location	Downstream paper and anticipated use
MePMoS / NHT / HAL / BFFT hierarchy	Theorems 48.2, 50.2, etc.	GG-A8-QM: explicit named hierarchy carrying the action-conditioned IH observability and the runtime cascade of any persistent quantum-information system.
SeOS persistence criterion + QCI completeness index	Theorem 58.1 and Definition 59.1	GG-A10-Cosmos: persistence class for the cosmological boundary-projection of living open systems.
OH+IH=EIC (definitional short-hand $\text{EIC} := \mathcal{C}_{\text{loop}}(\text{OH}, \text{IH})$)	Section 64	GG-A8-QM: holographic-observation framework for the action-conditioned IH; GG-A10-Cosmos: cosmological-scale outer hologram.
Pagoda hierarchical factorization theorem (five-floor variant)	Theorem 66.2	GG-A9-SM: particle floor; GG-A10-Cosmos: cosmos floor; future papers on chemistry / life / mind floors. (Open problem O7 records that “five” is the historically natural variant, not yet uniquely forced by MSMF.)
Conditional inevitability theorem	Theorem 66.4	GG-A10-Cosmos: anthropic-style boundary-cosmology argument for the inevitability of life and mind in a sufficiently rich driven universe. (Open problem O8 records the unmeasured Borel–Cantelli hypotheses.)

The two maps together establish the position of Qi-4 in the GG-Theory program: the constitutional content flows downward from A1, the bulk-side technical content flows in from A2 and A3, the boundary-projection structure flows in from A5, the geometric pillar template flows in from A6, and the four Qi-pillars + SeOS/EIC then flow outward to every downstream paper that addresses the 4D-physics, cosmology, or predictions branches of the program. When any companion paper is later revised to a new version, the present section should be re-walked against the revised paper’s actual theorem numbering before the present paper is itself republished; the audit trail above is the working record of the current numbering snapshot.

74 The Ten Structural Innovations of the Qi-4 Runtime

The contribution of GG-A7-Qi4 is to carry the open-system boundary law of GG-A5-GDOS into a complete information-theoretic runtime — four pillars, an internal-measurement construction, and a cross-scale hierarchy — without introducing any new continuous parameter beyond the imported ledger. We collect here the ten structural innovations that runtime delivers, each carrying its cohort-global identifier (§74.1).

Innovation 1 (I-Qi4-1): Open systems as fundamental. The runtime takes open, environment-coupled systems as the primitive object rather than as corrections to ideal closed systems — the boundary law, not the isolated Hamiltonian, is the starting point.

Innovation 2 (I-Qi4-2): Observation as a constitutional projection. Measurement is treated as the lawful bulk-to-boundary OGGEF projection from the outset (Theorem T-Qi4-3.3), not as a late interpretive patch added to a closed theory.

Innovation 3 (I-Qi4-3): Information as a gauge current. Information is given physical, gauge-current status through the memory-bearing $U(1)_{Qi}$ sector and the Arisaka ledger $Q_B = (2/5) E_{Ryd}/K$, rather than remaining an informal metaphor.

Innovation 4 (I-Qi4-4): One runtime vocabulary from particles to society. The same seven-layer grammar — packet, ledger, selector, diagnostic, architecture, persistence, and holographic observation — is reused unchanged across chemistry, life, brains, and collective systems.

Innovation 5 (I-Qi4-5): The four pillars as principal theorems. The informational pillars Qi-QUA, Qi-EP, Qi-CON, and Qi-HAL are stated as principal theorems with the implication chain $Qi-QUA \Rightarrow Qi-EP \Rightarrow Qi-CON \Rightarrow Qi-HAL$, each forcing the next rather than standing as independent assumptions.

Innovation 6 (I-Qi4-6): The Arisaka four-current lock. Qi-QUA ties the four boundary currents through one energy–information conversion factor, $J_E/Q_B = J_e = J_B = J_{Qi}$ (the rank-one inflow theorem T-Qi4-9.3), so that the informational current is not free but locked to the energetic, electric, and baryonic channels.

Innovation 7 (I-Qi4-7): Three selector gauges unified. The action-extremum, current-statistics, and inferential readings of Qi-EP (PLT/MEPP/FEP) are shown to be three gauges of one projected functional (the common-origin theorem T-Qi4-10.6), not three competing principles.

Innovation 8 (I-Qi4-8): The two-knee conformal signature. Qi-CON predicts a specific two-knee shape for $1/f$ spectra — exactly two scale-breaks bracketing a finite power-law band, a structural signature distinct from a single unbounded slope (the two-knee spectral theorem T-Qi4-11.7), with the knee positions set by the system’s slowest and fastest retained rates — a sharp, multiscale falsifier rather than a generic power law.

Innovation 9 (I-Qi4-9): Internal measurement and the late placement of consciousness. The Outer-Hologram / Inner-Hologram / Evolving-Inner-Cosmos construction (OH+IH=EIC, universality theorem T-Qi4-15.14) places conscious content late in the derivation, as the internal measurement of selected IH structure, rather than early as an unexplained primitive.

Innovation 10 (I-Qi4-10): The internalized GRIP-loop and the cross-scale pagoda. Qi-HAL discovers that a memory-bearing system carries its own recoverable GRIP-Out (the internalized regime switch), and the five-story pagoda factorization (T-Qi4-16.2) exhibits the same runtime from particle to mind as a single nested hierarchy.

74.1 Global naming convention for cross-cohort citation

Beginning with [Arizaka \(2026, A2-GG6\)](#) and extended across the cohort, each load-bearing item of the GG-Theory program carries a globally unique identifier of the form *Prefix-Suffix-N*, where the Prefix names the item type and the Suffix names the paper:

- **T-Suffix-N** for the *N*th Theorem (GG-A7-Qi4's theorems are section-numbered, so its global ids carry the section number, e.g. T-Qi4-9.3 is the rank-one inflow theorem of Qi-QUA and T-Qi4-15.14 the EIC universality theorem).
- **L-Suffix-N**, **C-Suffix-N**, **D-Suffix-N**, **R-Suffix-N** for the *N*th Lemma, Corollary, Definition, and Remark.
- **P-Suffix-N** for the *N*th Prediction.
- **O-Suffix-N** for the *N*th Objection (e.g., O-Qi4-1 is the first Common Objection of the present paper).
- **F-Suffix-N** for the *N*th Falsifier (e.g., F-Qi4-1 is the current-locking falsifier of the empirical program).
- **I-Suffix-N** for the *N*th Innovation (e.g., I-Qi4-1 is the open-systems-as-fundamental innovation above).
- **PS-Suffix-N** for the *N*th Problem-Solved / open-problem entry (§75.2).

The canonical Suffix list for the twenty-two cohort papers is: *Constitution* (A1), *GG6* (A2), *BI6* (A3), *Euclid* (A4), *GDOS* (A5), *GRIP* (A6), *Qi4* (A7, this paper), *QM* (A8), *SM* (A9), *Cosmos* (A10), *Origin* (A11), together with the phenomenology preprints *KK* (P1), *N3* (P2), *Electron* (P3), *Proton* (P6), the foundations preprints *Real* (P4) and *CPT* (P5), and the consequence preprints *Kaxion* (C1), *Hubble* (C2), *BH* (C3), and the white papers *Symmetry* (W1) and *Holonomy* (W2). The four axioms (AX-1... AX-4) and seven postulates (POS-1... POS-7) of [Arizaka \(2026, A1-Constitution\)](#) keep their existing cohort-wide names; the global identifier is the same item viewed from outside its home paper.

75 Open Problems, Limitations, and Falsifiability

Assessment has a second, harder side. After significance and comparison, the manuscript must say where the framework could fail and how it should be tested. Section 75 therefore isolates limitations, falsifiability, and the empirical program into a separate assessment layer.

75.1 Why assessment must be split from significance

A landmark synthesis must not only say why it matters; it must also say where it is incomplete and how it could fail. For that reason the revised architecture separates assessment into two layers. The previous section addressed significance, uniqueness, and relation to existing programs. The present section addresses the harder question: what remains unfinished, and what concrete empirical failures would count against the framework?

75.2 Open problems and limitations of the derivational stage

A serious discussion must state open problems and limitations clearly. Following the L6 audit convention adopted across the GG-Theory program, each open problem O_k below is triaged into one of three status categories:

- *Resolved.* The question has a complete answer in a named companion paper, with explicit citation.
- *Partially resolved.* The central claim is settled in the named companion paper, with a specific sub-task still pending in that companion paper.
- *Genuinely open.* No companion paper currently carries the answer. The body identifies the candidate mechanisms or research directions.

Listing as “open” what is in fact resolved in a companion paper would be a methodological failure of the constitutional program; the triage below records the actual status of each question. A closing summary at the end of the section gives the status counts.

(O1) (PS-Qi4-1) Universal numerical simulation package spanning all pagoda floors.

Status: Genuinely open. The theory is architectural before it is numerically exhaustive. The present paper gives theorem-level derivations, normal forms, and reduction criteria, but it does not yet supply one universal simulation package spanning all five floors of the pagoda. Domain-specific numerical models will still be needed. A unified simulation toolkit is a substantial software-engineering project deferred to future work.

(O2) (PS-Qi4-2) Measurement protocols for $U(1)_{Qi}$ at chemistry, life, and brain scales.

Status: Genuinely open (experimental, awaiting protocol development). The empirical realization of $U(1)_{Qi}$ at the chemistry, life, and brain scales is at present a structural and bookkeeping claim rather than a finished measurement protocol. The framework says what must be tracked as informational current; it does not yet provide a single off-the-shelf instrument family for every domain. Direct detection of the Kaxion-like $U(1)_{Qi}$ signature is the high-energy companion problem and is tracked in [Arizaka \(2026, C1-Kaxion\)](#).

(O3) (PS-Qi4-3) Full phenomenology of consciousness. *Status: Genuinely open.*

The consciousness discussion is deliberately disciplined but incomplete. The framework identifies a runtime condition for conscious content — internally measured selected IH structure — without claiming to have solved the entire phenomenological problem. The criterion the framework

supplies is one of *access* consciousness — the reportability and global availability of selected content, in the sense of Block (1995); the *phenomenal* hard problem of subjective experience is recorded here as a structural placeholder for a future companion paper.

- (O4) **(PS-Qi4-4) Individual-environment realization of the pagoda hierarchy.** *Status: Partially resolved.* The hierarchy theorem is proved in the ensemble sense earlier in the present paper: averaged over admissible environments, the higher floors are inevitable. The individual-environment realization is conditional on accessible viable windows having positive measure and on sufficient driven exploration. No companion paper currently addresses the local-environment refinement, though some sub-questions are tracked in Arizaka (2026, A10-Cosmos).
- (O5) **(PS-Qi4-5) OGN-5 selection rule for the Arizaka constant Q_B .** *Status: Genuinely open.* The Arizaka constant $Q_B \approx 0.1555 \text{ eV/qubit}$ is defined by $Q_B = (2/5) E_{\text{Ryd}}/K$ with $K = 35$ from the OGN-5 ledger. The assignment “per-step quantum = E_{Ryd}/K ” is currently justified only by minimality intuition; a derivation showing E_{Ryd}/K (rather than $E_{\text{Ryd}}/J = E_{\text{Ryd}}/7$ or another combination of the OGN-5 integers) is the unique per-step quantum under MSMF would close this question. The earlier factor-of-two gap to ATP hydrolysis is resolved: one phosphoanhydride bond carries two locked quanta, so $2 N_A Q_B = 30.0 \text{ kJ/mol}$ matches the measured standard free energy $-30.5 \pm 0.5 \text{ kJ/mol}$ to 1.6% (Part III). What remains genuinely open here is the selection-rule proof of the functional form E_{Ryd}/K , not the bioenergetic magnitude.
- (O6) **(PS-Qi4-6) Uniqueness of the MePMoS-NHT-HAL-BFFT factorization in the GRIP cascade.** *Status: Genuinely open.* Theorem 53.1 establishes that the named five-layer hierarchy is a *sufficient* boundary realization of the GRIP cascade on the family-projector + memory moduli. Uniqueness — the claim that no other 3-stage, 5-stage, or category-equivalent factorization satisfies the cascade — requires a separate no-go argument that is currently absent.
- (O7) **(PS-Qi4-7) Derivation of the pagoda floor count = 5 from MSMF.** *Status: Genuinely open.* The five-story pagoda (particle $\rightarrow$ cosmos $\rightarrow$ chemistry $\rightarrow$ life $\rightarrow$ mind) is asserted as the intended interpretation of the hierarchical factorization theorem. Whether “five” is forced by AX-4 MSMF (minimality across closure-defect and port-novelty inequalities) or whether other floor counts satisfy the same admissibility criteria is currently open. A four-floor or six-floor variant would satisfy the existence theorem on its face.
- (O8) **(PS-Qi4-8) Quantitative bounds on the conditional-inevitability hypotheses.** *Status: Genuinely open.* Theorem 66.4 reduces the inevitability of life and mind to a Borel–Cantelli application conditional on (i) finite invariant measure on the driven accessible region, (ii) positive measure of the viable subsets W_{life} , W_{mind} , and (iii) per-visit stabilization probabilities bounded below. The unconditional inevitability claim depends on independent bounds for $\mu(W_{\text{life}})$ from chemistry-floor analysis and p_{life} from SeOS persistence statistics; these are not provided in the present paper.
- (O9) **(PS-Qi4-9) Three-component Qi-4’ framework.** *Status: Genuinely open.* The L12 framing of $U(1)_{Q_i}$ admissibility (above) acknowledges that a hypothetical three-component theory

$\text{Spin}(1,5) \times \text{SU}(5) \times \text{U}(1)_B$ would require an entirely new Qi-4' framework: a four-current Arisaka equation reducing to three currents, a memory carrier other than a gauged $U(1)$, and a reformulated EIC construction. Whether such a Qi-4' is consistent at all, and what its empirical predictions would be, is recorded here for future work.

- (O10) **(PS-Qi4-10) The converse of the GGA characterization.** *Status: Partially resolved.* Theorem 62.1 proves that the ground attractor of an admissible stratum carries all four certificates. The converse — that no other critical point of the stratum carries all four — is settled on admissible strata by GG-A6-GRIP's per-stratum corollary (Arisaka, 2026, A6-GRIP): there is exactly one critical point of index 0, every other has a direction of descent, and the halting certificate fails there (Remark 62.3). What remains open is the converse on strata whose admissibility is open — the Wilson-line stratum first — where a landscape can carry two floors of equal depth that Qi-HAL cannot separate, and it is not known whether the ledger, variational, or scale certificate can. The sub-task pending is therefore GG-A6-GRIP's: establish or refute admissibility on those strata. The openness is inherited from that question, not produced by the pillars.

These open problems are not defects hidden in small print. They are the honest boundary of the present derivational stage, and the triage above records where each lives in the GG-Theory program.

75.3 Falsifiability and empirical program

The framework would be weak if it only re-described existing phenomena after the fact. Its scientific value depends on linked falsifiable predictions. The most important point is that the predictions are *joint*. Many individual signatures, such as power laws or prediction errors, can appear in other frameworks too. The present framework predicts that specific families of signatures should co-occur because they descend from one runtime.

Define a system-level consistency vector

$$\Sigma_S := (C_{\text{lock}}, C_{\text{sel}}, C_{\text{scale}}, C_{\text{mem}}, C_{\text{holo}})_S, \quad (75.1)$$

where the components measure, respectively: evidence for current locking (Qi-QUA); evidence for variational behavior (Qi-EP); evidence for scale-covariant diagnostics (Qi-CON); and two faces of Qi-HAL — evidence for the memory architecture (MePMoS, NHT, HAL, BFFT) and evidence for the internality layer ($\text{OH} + \text{IH} = \text{EIC}$). The last two are scored separately because a system can implement storage and timing richly while closing no model-of-the-environment loop at all, and that dissociation is exactly what the vector is for. The framework predicts not just isolated positive components, but positive *coherence* across the vector in systems claimed to instantiate the full runtime.

A simple scalar summary is

$$\Xi_S := \sum_{a=1}^5 w_a C_a(S), \quad w_a > 0, \quad \sum_a w_a = 1. \quad (75.2)$$

The exact operationalization of C_a depends on domain, but the logic is universal: a true full-runtime system should score coherently across all five categories.

Table 54 lists representative tests.

Table 54: Representative falsifiability program for the integrated framework. Each row is one predicted signature, with a test idea and what a failure would mean. The third column is the one that gives the table its value: a framework this broad can absorb a surprising amount of data unless it says in advance which failures are fatal and which are merely inconvenient, and this column distinguishes the two case by case.

Predicted signature	Test idea	What failure would mean
Current-locking constraints	Estimate event-level energy/charge/information budgets in repeated micro-events; test whether one locked ledger explains them.	Persistent inability to define a coherent event ledger would challenge Qi-QUA.
Selector behavior	Perturb pathway options and test whether realized trajectories follow a common PLA/PLT/MEPP/FEPstyle selector.	If no selector family fits across domains, Qi-EP loses force.
Two-knee scale signatures	Measure spectra, burst statistics, and fractal indicators under controlled conditions.	Absence of multiscale diagnostic structure would weaken Qi-CON.
Exploration-to-commitment lag	Time-resolve candidate-state exploration, branch selection, and stable report or action.	If selection/stabilization staging is absent, the DGI $\rightarrow$ GSSB $\rightarrow$ GA pipeline is incomplete.
Motion-conditioned model quality	Restrict action while preserving passive input and test whether internal model quality degrades.	Failure would challenge the action-conditioned observability theorem behind IH.

The paper is therefore falsifiable in a strong sense. It is not protected by saying “something complicated happened.” It predicts specific linked structures: packet stability, ledger accounting, selector behavior, scale diagnostics, memory architecture, and holographic observation semantics. If those structures fail jointly, the theory fails.

Experimental open-problem list (L6 triage). We catalog the five principal experimental directions of Table 54 as labeled open-problems E1–E5, each with the L6 status header. All five are presently *Genuinely open* in the experimental sense — the framework provides the prediction, the laboratory verdict is the next decade of measurements.

- (E1) **(F-Qi4-1) Current-locking constraints.** *Status: Genuinely open* (awaiting empirical data). Event-level energy/charge/information budgets in repeated micro-events should be consistent with one locked ledger per Qi-QUA. Falsification: persistent inability to define a coherent event ledger across the four boundary currents J_e , J_B , J_{Qi} , J_E .

- (E2) **(F-Qi4-2) Selector behavior across domains.** *Status: Genuinely open.* Perturb pathway options and test whether realized trajectories follow a common PLA/PLT/MEPP/FEP-style selector. Falsification: no selector family fits across domains.
- (E3) **(F-Qi4-3) Two-knee scale signatures.** *Status: Genuinely open* (predicted by T-CON.2 of Qi-CON, awaiting controlled spectral measurements). Specific predicted shape of $1/f$ noise with two co-occurring scale-breaking knees bracketing a finite power-law band (the structural signature, distinct from a single unbounded $1/f$ slope), with the knee positions set by the system’s slowest and fastest retained rates rather than by universal constants.
- (E4) **(F-Qi4-4) Exploration-to-commitment lag (DGI→GSSB→GA timing).** *Status: Genuinely open.* Time-resolve candidate-state exploration, branch selection, and stable report or action. Falsification: selection/stabilization staging is absent or has the wrong temporal structure.
- (E5) **(F-Qi4-5) Motion-conditioned model quality (action-conditioned IH observability).** *Status: Genuinely open.* Restrict action while preserving passive input and test whether internal model quality degrades. The qualitative active-versus-passive asymmetry is already established ([Held & Hein \(1963\)](#); cf. the reafference principle of [von Holst & Mittelstaedt \(1950\)](#)); the sharper, still-open prediction is the *quantitative* dependence of model error on the observability-Gramian deficit. Falsification: failure of that quantitative dependence would challenge the action-conditioned observability theorem behind IH.

Triage summary. Of the fifteen open problems cataloged above (O1–O10 architectural and theoretical, E1–E5 experimental), *zero* are Resolved in a named companion paper at this writing; *two* are *Partially resolved* — O4 (individual-environment realization of the pagoda), with the central claim settled in the ensemble sense in the present paper but with sub-tasks pending in [Arisaka \(2026, A10-Cosmos\)](#), and O10 (the converse of the GGA characterization), settled on admissible strata by GG-A6-GRIP’s corollary and pending on the strata whose admissibility that paper leaves open; the remaining *thirteen* (O1, O2, O3, O5–O9, E1–E5) are *Genuinely open* and constitute the live agenda for the next phase of the Qi-4 program. The open problems O5–O10 record the uniqueness and selection-rule gaps that a hostile reading of the construction exposes: the OGN-5 selection rule for Q_B (O5), the uniqueness of the Qi-HAL five-layer factorization (O6), the derivation of the pagoda floor count from MSMF (O7), the quantitative bounds on the Borel–Cantelli inevitability hypotheses (O8), the construction of a three-component Qi-4’ framework (O9), and the converse of the characterization theorem off the admissible strata (O10). Indeed, the proportion of Genuinely-open entries is the natural signature of a paper at the derivational frontier: most of the constitutional content is settled; the uniqueness and selection-rule refinements — and the laboratory verdict — are the work of the next decade.

76 What a framework is entitled to claim

There comes a point in a construction of this size when the useful thing to do is stop building and take inventory, and this Part has been that point.

Four things are now on the record and are better said plainly than hedged. The framework is architectural before it is numerical: it says what kinds of quantity there are, how they constrain one another, and what would count as a violation, and for most domains it does not yet say what the numbers are — that work needs simulation and calibration this paper has not done. Its information-carrying gauge structure, outside the one sector where it was derived, is at present a bookkeeping claim: well defined, internally consistent, and still short of the experimental proxy that would make it a measurement rather than a convention. Against those two admissions stands the reason the construction is worth arguing with at all — what it predicts is not a handful of separate effects but a web of linked ones, so it can be caught out by a disagreement between two of its own predictions, with nobody obliged to take the framework's word for anything. And from all three together follows the fourth: its value is hostage to a coordinated program of measurement that does not yet exist, and cannot be rescued by reinterpretation if that program returns the wrong answer.

That last sentence is the one worth dwelling on, because the alternative is so much more comfortable and so common. A framework broad enough to speak about particles and cells and institutions is under continuous temptation to make itself unfalsifiable by degrees — never in one decision, always in small ones, each individually reasonable. A parameter is left free because the domain is complicated. A disagreement becomes a matter of the regime not having been the one intended. A retrodiction is counted as a success. None of these is a lie, and their accumulation is fatal, because what is left at the end is a vocabulary that can describe any outcome and therefore expects none. The discipline this Part tries to impose is the refusal of each of those moves one at a time, and the place where the refusal shows is the open-problem list: it is long, most of its entries are genuinely open rather than pending, and the ones that would hurt most are named in the same voice as the ones that would not.

There is a second reason to keep the list long and honest, and it has nothing to do with modesty. A theory's open problems are the only part of it that tells you where to go next. A construction with no acknowledged gaps offers a reader nothing to do but agree or leave. A construction that says exactly which joint is glued, which step is conditional, and which number was fitted rather than derived is handing over a map with the unexplored regions marked — and those regions, not the settled interior, are where the next result will come from. On this reading the ratio of open to closed in a young framework is not an embarrassment to be minimized in the abstract. It is a statement about how much of the thing is still alive.

What can be said, then, at the close of the assessment. Not that the account is established — it is not, and the preceding sections have been specific about which parts are proved, which are conditional on named hypotheses, and which are organizational readings that could be replaced tomorrow without the structure noticing. What can be said is that the account has been put into a shape where being wrong is possible, and where the ways of being wrong have been enumerated by the author rather than left for somebody else to discover. That is a smaller thing to have achieved

than a theory of nature. It is also the only condition under which the larger claim, if it is ever earned, would mean anything at all. The final Part sets out the whole program in the order in which it was built, so that a reader who wants to disagree can see precisely where to begin.

Part XII

Conclusions

Synthesis of the Qi-4 runtime, the falsifiable program, and the constitutional reading of open-system physics

77 Conclusions

This paper asked a single question: once a lawful six-dimensional bulk has been established and a lawful four-dimensional boundary has been derived from it, what universal runtime must govern any open system that lives on that boundary? The answer developed across the paper is that the runtime is not a loose metaphor for openness. It is a definite mathematical structure—the GDOS program with the Qi-4 Pillars—that can be stated, tested, and falsified.

The derivation proceeded in a deliberate order, and that order is itself part of the result. One cannot meaningfully lock currents (Qi-QUA) before declaring the packet that carries them (SIP) and the ledger that budgets them (GFF). One cannot meaningfully diagnose multiscale signatures (Qi-CON) before establishing the variational grammar that generates trajectories (Qi-EP). One cannot meaningfully discuss consciousness (EIC) before establishing what measurement, persistence, and internal modeling mean in the runtime. This dependency-ordered architecture is therefore not an editorial convenience but a logical dependency chain: each layer requires the one below it and enables the one above it.

77.1 The synthesis theorem

The full derivation chain can now be stated as a single theorem.

Theorem 77.1 (GDOS synthesis theorem (T-Qi4-22.1)). *Assume the constitutional results of Papers I–IV-A: the GG-Constitution axiom system, the GG-6 bulk with BI-6 and OGN-5, and the OGGEF-derived observer boundary (the imported GRIP-cascade closure remaining conditional on M11 in its spectral-flow form, M10 resolved externally and the scope refinements closed at class level). Then any natural open system that admits an accessible/hidden split, persistent reduced records, and coarse-grained current closure is described by the following nested chain:*

- (i) *DGI+ GSSB+ GA as the universal runtime triad;*
- (ii) *GFF as the universal scalar ledger of energy, entropy, and free energy;*
- (iii) *SIP as the canonical packet that instantiates the ledger in a concrete open system;*
- (iv) *Qi-QUA as the locked current ledger, giving the Arisaka Ratio, Arisaka Equation, and Arisaka Constant;*

- (v) *Qi-EP as the extremal grammar whose boundary gauges include PLA, PLT, MEPP, and FEP;*
- (vi) *Qi-CON as the scale-covariant sector producing heavy tails, fractality, and double-knee $1/f$ statistics;*
- (vii) *Qi-HAL as the minimal architectural sector for persistent adaptive information processing;*
- (viii) *SeOS and QCI as the persistence/completeness layer of self-maintaining open systems;*
- (ix) *OH, IH, and EIC as the dual-hologram closure for systems capable of internal reconstruction;*
- (x) *the pagoda hierarchy as the graded map from particle to cosmos, chemistry, life, and mind.*

This chain is stable under lawful coarse-graining in the sense that lower stories can generate higher stories through story-lift reductions, while higher stories can constrain lower stories through admissible control couplings.

Proof. We list, for each clause (i)–(x), the specific theorem of the present paper (or its companions) that supplies it, together with any inherited hypotheses. The synthesis theorem then holds under the union of these hypotheses.

- (i) DGI + GSSB + GA as the universal runtime triad: derived in Section 7, Theorem 7.3; inherited from GG-A6-GRIP (Arisaka, 2026, A6-GRIP) Theorem T4 (GRIP cascade). Hypothesis: an admissible bulk action satisfying OGGE.
- (ii) GFF as the universal scalar ledger: derived in Section 8, Theorem 8.2; inherited from GG-A6-GRIP Theorem T3. Hypothesis: a smooth boundary state manifold with metric g and gauge connection $\mathcal{A}$.
- (iii) SIP as the canonical packet: derived in Section 9, Theorem 9.2. Hypothesis: the system admits a finite-dimensional accessible state space with a declared observable algebra.
- (iv) Qi-QUA as the locked current ledger: derived in Part III, Theorems 12.2 and 15.3. Hypothesis: the four-component bulk gauge group $\text{Spin}(1,5) \times \text{SU}(5) \times \text{U}(1)_B \times \text{U}(1)_{Q_i}$ and the BI-6 closure flowing through the rank-one inflow channel.
- (v) Qi-EP as the extremal grammar: derived in Part IV, Theorems 27.1, 28.1, 29.1, and 32.1. Hypothesis: tree-level / large-deviation reduction of the influence functional (the boundary saddle-point assumption).
- (vi) Qi-CON as the scale-covariant sector: derived in Part V, Theorems 39.1 and 42.1. Hypotheses: the admissibility window $0 < \alpha \leq 2$ on the fractional exponent and the soft-mode density assumption $A_X(\lambda)\rho(\lambda) \approx A_0\lambda^{\gamma-1}$.
- (vii) Qi-HAL as the minimal architectural sector: derived in Part VI, Theorem 53.1. Hypothesis: a memory-bearing open system with active $\text{U}(1)_{Q_i}$ coupling.
- (viii) SeOS and QCI as the persistence/completeness layer: derived in Part VII, Theorem 58.1 and the QCI construction therein. Hypothesis: the maintenance-margin inequality of Part VII.

- (ix) OH, IH, EIC as the dual-hologram closure: derived in Section 64, Theorems 64.5, 64.9, and 64.14. Hypothesis: continued re-measurement and a sufficiently rich internal-model class.
- (x) The pagoda hierarchy: derived in Section 66, Theorem 66.2. Hypothesis: floor-by-floor admissibility (closure defect, lifetime separation, port-rank enlargement); the empirical termination at floor 5 is conditional on Open Problem O7.

The story-lift property between adjacent floors is supplied by Theorem 66.2 (clause (x)). The admissibility-after-coarse-graining property is supplied by Theorem 67.2 of Section 4 (the universality theorem of Part I). The conditional-inevitability statement at floor 5 is supplied by Theorem 66.4 of Section 66 under the four hypotheses identified in Open Problem O8.

Therefore the whole chain follows under the union of all hypotheses listed above. ■

Remark 77.2 (On the conditional character of the synthesis). Several of the clauses listed in the proof above inherit hypotheses that the present paper itself flags as Genuinely-open in the Open Problems triage of Section 75. Specifically:

- clause (vi) (Qi-CON) depends on the soft-mode density assumption and the admissibility window $0 < \alpha \leq 2$;
- clause (vii) (Qi-HAL) depends on the active $U(1)_{Q_i}$ coupling, which is constitutionally optional and empirically preferred (see the L12 framing in Section 75);
- clause (viii) (SeOS, QCI) depends on the calibration protocol for the packet-distance weights and the completeness-score thresholds (Open Problem O2);
- clause (x) (the pagoda hierarchy) depends on floor-by-floor admissibility, conditional on Open Problem O7.

Theorem 77.1 is therefore best read as a *conditional* synthesis statement: the chain holds under the union of all listed hypotheses, four of which (highlighted above) are themselves the subjects of explicit open problems in the present paper. The empirical mission of the manuscript is to make those open problems concrete enough that a falsification at any single hypothesis collapses the chain at a named clause, with the rest of the synthesis remaining intact.

Theorem 77.1 is the main deliverable of the paper. It does not mean that all detailed science is finished. It means that the constitutional and runtime scaffolding for a unified open-system program is now explicit, auditable, and refutable.

77.2 Three principal conclusions

Three conclusions deserve emphasis because they go beyond the technical content of any single section.

First: the open-system world has a definite runtime geometry. The strongest single claim of the paper is that a universal science of open systems requires one common runtime that begins with lawful geometry and ends with lawful observation, memory, persistence, and hierarchical emergence. Before the present work, these capacities were addressed by separate traditions—quantum measurement theory, nonequilibrium thermodynamics, dynamical systems, active inference, and consciousness science—each with its own foundational language. After the present work, they are shown to be successive layers of one derivational chain. The differences between domains appear in the SIP data, not in the constitutional law.

Second: life and mind belong inside physics, not beside it. Life is the floor of the pagoda at which maintenance-positive open packets become stable and evolvable under the SeOS criterion. Mind is the floor at which the inner hologram becomes sufficiently rich and self-accessible to support internal measurement of selected content. Neither requires a separate ontological category. Neither is merely reduced to the floor below it. The pagoda hierarchy explicitly says that new stories emerge by new slow variables, new selection rules, and new attractor families. Chemistry is not a useless paraphrase of particle physics. Life is not a useless paraphrase of chemistry. Mind is not a useless paraphrase of life. What the present theory denies is not level-specific reality; it denies level-specific *lawlessness*. Every floor must obey the constitutional runtime, and every floor genuinely extends what the floors below it can do.

Third: the framework is reflexive. Because the framework applies to distributed memory-bearing systems, scientific communities become examples of the runtime they study. Experimental data are OH: lawful projections from a larger world to a registered boundary record. Theories are IH: internally constructed models that attempt to reconstruct the structure behind those records. The history of science is an EIC refinement process, carried out collectively, in which the mismatch between outer data and inner model is gradually reduced. This reflexive feature is not an embarrassment. It is an expected consequence of a genuine theory of observation and model-building, and it provides a natural account of why science works, why it is difficult, and why it converges.

77.3 Outlook and the next generation of papers

The present paper is broad by design: it had to be, because the runtime must be stated once and completely before domain-specific work can begin on firm constitutional ground. The next generation of papers will be narrow by design, each testing one sector of the runtime against one body of data.

Chemistry and the origin of life. Papers centered on GFF and Qi-QUA at the molecular scale will attempt to verify the Arisaka Equation against known catalytic cycles (ATP synthase, proton-coupled electron transfer, kinetic proofreading), derive the SeOS transition from protocells to minimal life, and predict testable current-locking signatures in single-molecule experiments.

Neuroscience and consciousness. Papers centered on Qi-HAL and $OH + IH = EIC$ will develop quantitative models of hippocampal phase coding as NHT, cortical attractor landscapes as HAL, multiscale oscillatory coordination as BFFT, and the staged delay from stimulus to reportable conscious content as the $DGI \rightarrow GSSB \rightarrow GA$ pipeline applied to neural populations. Predictions about exploration-to-commitment timing, motion-conditioned model quality, and inner-hologram degradation under specific lesions will be stated and compared with existing neurophysiological data.

Artificial and engineered systems. Papers testing whether engineered systems (autonomous robots, large-scale AI architectures, synthetic biological circuits) can be made genuinely GDOS-admissible will provide a controlled environment in which the runtime can be verified or falsified without the complexity of natural systems.

Collective and social systems. Papers on distributed social, economic, and institutional open systems will formalize collective inner cosmos formation, test whether institutional persistence obeys SeOS-type criteria, and ask whether the consistency vector Σ_S can be measured for organizations and scientific communities.

Quantitative sharpening. Across all domains, the deeper task is cumulative calibration: turning the architectural framework into numerical predictions with stated error budgets. Every sector of the current paper—from the Arisaka Constant to the two-knee spectral template to the QCI scoring protocol—can be made more explicit by model class, data type, and measurement instrument. This sharpening is now a matter of development rather than of constitutional incompleteness. The program has crossed the threshold from scattered motifs to a connected derivational scaffold.

77.4 Closing reflections: four pillars, one object

Epilogue. The remainder of this subsection steps out of derivation mode and into reflective commentary on what the framework attempts and what it promises. Readers expecting a technical close may skip directly to the Appendices.

We conclude with a broader reflection on what the present paper attempts to do and what it hopes, in this volume, to have partially accomplished. The technical content of the four Qi-4 pillars is in place; the philosophical content of those pillars deserves its own paragraph.

In 1944, in a slim book that has never quite gone out of print, Erwin Schrödinger asked a question that physics still has not settled: *what is life?* His proposed answer was that life is what happens when an open system extracts a stream of order from its environment and uses that stream to maintain itself against the tendency of entropy to dissolve it. He believed, but could not prove, that this would one day be a piece of physics — not a metaphor borrowed from physics, but physics in the same sense that gravity and electromagnetism are physics. Three quarters of a century later, the question still sits at the edge of every textbook of biophysics, every paper on the origins of life, every introductory remark in a neuroscience lecture. It has not been answered because the language has not been written.

The present paper proposes a language. It identifies a single quantity — the current of information — and gives it the same standing as the current of electric charge that runs through a copper wire, or the current of energy that crosses the surface of the Sun. Charge is conserved because the world has a $U(1)$ symmetry of electromagnetism; energy is conserved because the world is invariant under time translation; the proposal here is that information is conserved, along its appropriate boundary and on appropriate time-scales, because the world is invariant under a further $U(1)$ symmetry that lives on the six-dimensional bulk and projects to every open system as a fourth-current ledger. The conversion factor between the information current and the energy current of the same open system is a single small number, of size 0.1555 electron-volts per qubit, that we call the Arisaka constant. This number is not a fit. One molecule of ATP carries two of these quanta, so that twice the Arisaka constant, $2Q_B \approx 30.0$ kJ/mol, is the standard free energy of hydrolysis of the ATP phosphoanhydride bond (measured -30.5 ± 0.5 kJ/mol, an agreement to 1.6%) — the molecular currency that every living cell on Earth uses to keep itself alive. That a constant derived from a six-dimensional geometry should match a number that every biochemist has been measuring since the 1930s is the empirical surprise around which the present paper is organized, and it is the single sharpest reason to take the framework seriously.

If this picture survives the laboratory tests of the coming decade, what will have changed? At the level of the physics textbook, the chapter on conservation laws will gain a fourth entry alongside the energy, charge, and momentum currents that have been on the page for a hundred years. At the level of biology, the open-system character of every cell will stop being a complication to be handled after the closed-system idealization has been written down, and will instead become the starting axiom — every system in nature is open, life is what happens when openness is high-quality enough to sustain a memory, and consciousness — in the *access* sense the framework can address — is what happens when the memory becomes rich enough to model the system's own outer environment from the inside, while the *phenomenal* hard problem is left explicitly open. At the level of physics, the long-standing puzzle of how to fit the observer into the picture will lose its strangeness; the observer is not a complication added afterwards but the very thing the framework was built to describe. At the level of the philosophy of science, the relationship between physics and the life sciences will no longer be one of disciplinary courtesy: the same equations, the same currents, the same conservation laws will be available to the chemist studying catalysis, the cell biologist studying mitochondria, the neuroscientist studying the hippocampus, and the cosmologist studying horizons.

The deeper philosophical content of the framework is what it says about the relationship between the physical world and the living systems within it. In the standard reading inherited from the late nineteenth century, the universe is a collection of fields obeying laws that exist independently of any observer; living systems are an accidental, late, locally complex phenomenon that the universe produces in a tiny corner of itself, and that physics studies separately, if at all. In the Qi-4 reading, this division dissolves. The laws of physics and the living systems that obey them are not two ontological categories but two aspects of one foundational object — the six-dimensional bulk and the boundary on which all measurements occur. The bulk is not a remote stage on which the boundary world plays out; the boundary world is the place where the bulk geometry writes itself down in bits. Life is one specific way that the writing becomes legible. Mind is one specific way that the legible

writing folds back on itself and produces an internal model — an Evolving Inner Cosmos, in the operational sense of $\text{OH}+\text{IH}=\text{EIC}$ — that mirrors the outer environment well enough to predict, decide, and persist.

This is not a metaphor; it is the constitutional content of the operational/open extension of the Geometry-Gauge Equivalence Principle (AX-3 OGGE) and the four Qi-4 pillars that descend from it. Every measurement — the one a physicist makes with a particle detector at CERN, the one a neuron makes when it spikes in a sensory cortex, the one a cell makes when it commits to a fate decision based on the molecular concentrations in its environment, the one a scientific community makes when it ratifies or refutes a theory — is, in this reading, an OH event. Every internal model — the one the brain builds of the body's position in space, the one a cell builds of its homeostatic state, the one a community builds of the future of the climate it inhabits — is an IH object. The pair, taken together, is the Evolving Inner Cosmos: the way a living open system, embedded in a larger geometry it cannot fully access, navigates that geometry by mirroring it from the inside. In the older language of John Wheeler, the universe is an observer-participancy. In the language proposed here, this is no longer a slogan; it is the operational content of the four Qi-4 pillars, made falsifiable in four named laboratories.

For three and a half centuries, physics has been the science of what nature does when nobody is looking. The closed Lagrangian, the unitary wave function, the isolated universe — each was an attempt to step outside the world and write down its laws from nowhere. Each succeeded spectacularly within its own domain, and each failed at the moment a finite observer — a thermometer, a photographic plate, a brain — asked to be included in the picture. The axiomatic inversion proposed in these pages takes the finite observer not as a complication but as the starting point. Every system in nature is open; closure is the limit, not the law. The language in which physics writes its rules for the open-system world is the language of OGGE on the bulk and the four Qi-4 pillars on the boundary, with the energy–information Arisaka constant Q_B as the conversion factor between energy and information that the boundary obeys whenever it sustains a memory.

If this proposal is correct, the laboratory world of currents and free energies and spectral signatures and memory architectures will be recognized, in the coming century, as the boundary projection of one geometric structure that includes the observer in its formal syntax — not as a complication, not as an afterthought, not as a soft option to be added when convenient, but as a primitive object of the theory. The long-standing strangeness of the relationship between physics and the life sciences will, with appropriate humility, dissolve into a single grammar in which both speak. Schrödinger's question — what is life? — will not have a single one-sentence answer; but for the first time it will have an answer of the same kind that physics already gives to the question *what is light?* or *what is matter?* — an answer cast in the language of a conservation law, a symmetry group, and a coupling constant, falsifiable in any laboratory equipped to measure it.

If the proposal is wrong, the empirical record will say so within a decade. A measured discrepancy between $2Q_B$ and the standard free energy of ATP hydrolysis beyond the precision of modern calorimetry; a memory-bearing system whose $1/f$ spectrum does not exhibit the predicted two-knee shape; a neural population whose memory architecture decomposes outside the MeP-MoS/NHT/HAL/BFFT class; a controlled open system in which the three avatars of the extremal

principle (least action, maximum entropy production, free-energy minimization) disagree irreducibly — any one of these will collapse the upstream constitution and tell the next generation of researchers which clauses to keep and which to revise. The framework has been deliberately designed to expose itself to falsification at the axiomatic level; it is not safe by virtue of obscurity. Either the framework will succeed empirically, or its specific failures will sharpen the next attempt.

What the present paper proposes, in the end, is not a new theory of life, nor a new theory of mind, nor a new theory of information. It is a new contract between physics and the living world. Physics promises to write its rules down — not only for the closed systems of textbook tradition, but for the open systems that constitute every cell, every brain, every ecosystem, and every society we know. The living world, in turn, promises to falsify those rules if they ever lie. We hope, with appropriate humility, that the present paper contributes a useful step in that direction. Indeed, the next generation of physicists, chemists, biologists, and neuroscientists — who will inherit the empirical results that vindicate or falsify the framework — must be the final judges. After all, the question Schrödinger asked in 1944 was whether life requires a new kind of physics; the answer that emerges here is that life does not require new physics, but it does require physics to recognize that information has always been one of its first-class objects, and that the boundary on which we live is, in a precise and laboratory-testable sense, the place where the upstream geometry writes itself down in bits.

We close with three reflections in the spirit of the philosophical closing of [Arisaka \(2026, A6-GRIP\)](#). Each is one paragraph; together they record what the present paper believes it has accomplished.

The four pillars as one object. The four pillars Qi-QUA, Qi-EP, Qi-CON, Qi-HAL are not four constructions but four faces of one geometric object: a single composite connection on a single 6D principal bundle, projected onto a single boundary by a single CPTP channel, with the $U(1)_{Q_i}$ gauge sector carrying the memory. Each pillar names one face — the four-current ledger, the variational principle, the conformal-multiscale signatures, the memory architecture — and the four faces are mutually forcing. Once one face is fixed by the constitution and by the geometric pillars of [Arisaka \(2026, A6-GRIP\)](#), the other three are no longer free. This is the structural sense in which GG-Theory extends from particle physics through cosmos, chemistry, life, and mind in one program: the same geometric object is the upstream source, and the differences between domains appear only in which sectors are active.

The bargain with the laboratory. Every Qi-4 pillar carries an explicit, named, dimensional empirical falsifier. Qi-QUA forces the Arisaka equation $J_E/Q_B = J_e = J_B = J_{Q_i}$ with the energy–information conversion factor $Q_B \approx 0.1555 \text{ eV/qubit}$. Qi-EP forces three falsifiable laboratory regimes (action-extremum trajectories, current-statistics extremum, inferential-model-quality extremum) that should co-occur in any system claimed to instantiate the runtime. Qi-CON forces the two-knee shape of $1/f$ noise spectra. Qi-HAL forces the MePMoS/NHT/HAL/BFFT architecture in any memory-bearing $U(1)_{Q_i}$ system. Each falsifier is a prediction, not a hope. What makes a body of knowledge a science is not that it explains many things but that it forces a small number of things

to be exactly as they are; what makes that science honest is that the forcing carries with it the risk of being wrong. Qi-4 accepts that risk in four named places, across four laboratories that already exist.

The deeper claim. If the framework survives the experimental tests of the coming decade — ATP-hydrolysis bioenergetics matching $2Q_B$, two-knee $1/f$ signatures in controlled spectroscopy, the MePMoS/NHT/HAL architecture verified in neural populations, the conditional-inevitability bounds of the pagoda satisfied in cosmological measure — then the lesson of the four pillars will be that the laboratory world is the projection of a single geometric structure, and that the apparent multiplicity of informational currents, free-energy ledgers, scaling laws, and memory architectures is the multiplicity of the faces of one underlying object. If the framework does not survive, then this paper will remain as one technical record of what such a projection would have looked like, written down with sufficient rigor that the failure mode is itself diagnostic. Either way, the work has been worth doing: physics promises to write its rules down for the open-system world, and the world promises to falsify them if they ever lie.

Acknowledgments

This work was developed through extensive iterative collaboration with three large language models: ChatGPT (OpenAI) for exploratory drafting, narrative development, and broad-scope ideation; Perplexity for literature synthesis, citation verification, and recent experimental cross-referencing; and Claude (Anthropic) for structural organization, mathematical consistency checking, LaTeX typesetting, and final manuscript preparation. The sequential refinement of the GG-Theory framework (GG-Constitution / GG-6 / BI-6 / GDOS / Gi-4 Pillars / Qi-4 Pillars) across hundreds of conversation sessions with these three systems was essential to the coherence of the present paper.

The author thanks Aaron Blaisdell and Hirotaka Iwasaki for stimulating discussions on the program's interdisciplinary scope and for critical feedback on the accessibility of the presentation. The author also thanks hundreds of collaborators in the Elegant Mind Club over the past decade, including Javier Carmona, Elizabeth Milles, Mingda He, Syed Bashir Hydari, Brian Ta, Chris Dao, Umaima Afifa, Isabella Bestanoby, Caominh Le, and Natsuko Yamaguchi.

All scientific concepts (especially the foundation of the GG-Constitution, GGEP/GG-6/BI-6, and OGGEF/GDOS/Gi-4/Qi-4 Pillars), theoretical judgments, physical interpretations, and any remaining errors are the sole responsibility of the author.

Part XIII

Technical Appendices

Detailed proofs, cross-domain worked examples, and the complete abbreviation, symbol, and key-equation reference for the main text

The Appendices gather the technical machinery and the reference material that supports the main-text theorems of the Qi-4 pillars and of the OH+IH=EIC runtime completion. Roughly speaking, this is where the reader who wants to verify a specific derivation, run a specific worked example, or trace a specific symbol will spend time — while the reader who wants only the axiomatic content and the principal theorems of the main text can stop after Part XII and treat the back matter as a self-contained companion volume. Appendices A and B provide the abbreviations and the glossary of key symbols. Appendix C collects the key equations of the paper in a single ordered table. Appendix D unfolds the technical derivations for the GRIP pipeline. Appendices E–H supply the technical derivations for the four Qi-4 pillars (Qi-QUA, Qi-EP, Qi-CON, Qi-HAL). Appendix I supplies the technical details for SIP and GFF. Appendix J unfolds the technical derivations for the OH+IH=EIC construction. Appendix K gives the pagoda-hierarchy and cross-domain technical details. Appendix L collects the compact falsifiability matrix. Appendices M–O supply the chemistry, neural, and ecological worked-example panels for the Qi-4 runtime across scientific domains. Appendix P is the notation concordance and benchmark protocols. Finally, the inline bibliography with full Relevance annotations follows the appendices and closes the volume. Indeed, the fifteen Parts together with the sixteen appendices realize the Qi-4 runtime in full: Parts I–X build and apply the runtime, Part XI assesses it, Part XII concludes, and the back matter — the technical appendices, the common objections, and the references — supplies the infrastructure that makes the whole document self-checking.

A Summary of Abbreviations

Program-defined names are in sans-serif; standard physics abbreviations are plain Roman; the third column gives the primary Arisaka-GG paper and section.

Table 55: Every abbreviation used in this paper, with its meaning and the place it is introduced in full. Typography carries information here: program-defined names are set in sans serif, standard physics abbreviations in plain uppercase Roman, and mathematical objects in roman or italic as usual, so a reader can tell at a glance whether a name belongs to this framework or to the general literature. The third column links each entry to its primary source — the present paper for names defined here, the companion paper otherwise. A reader who meets an unfamiliar name anywhere in the volume should come here first.

Symbol	Meaning	Primary link
<i>Framework name and core program objects</i>		
GG-Theory	Grand Geometric Theory (also: Geometry-Gauge Equivalence Theory)	GG-A1-Constitution, Preface
GG-Constitution	The constitutional charter of the GG-Theory program: four axioms (AX-1–AX-4) and seven postulates (POS-1–POS-7)	GG-A1-Constitution, §2–§3
GG-6	Geometry-Gauge locking in six dimensions: the 6D bulk	GG-A2-GG6, §8 (T1-GG6)
GDOS	Geometric Dynamics of Open Systems: the boundary law implied by OGGEp	GG-A5-GDOS
BI-6	Green–Schwarz-modified Bianchi identity at 6D	GG-A3-BI6, §3
OGN-5	Occam-Gauge-Normalization 5-integer ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$	GG-A3-BI6, §7; T6-GG6
QGL	Quadruple Gauge Locking: the four-factor bulk gauge group $\text{Spin}(1, 5) \times \text{SU}(5) \times \text{U}(1)_B \times \text{U}(1)_{Qi}$	GG-A2-GG6, §10 (T3-GG6)
Gi-4	Geometric-Invariance four-pillar structure (TGG, CCI, GFF, GRIP)	GG-A6-GRIP
Qi-4	Quantum-Information four-pillar structure (Qi-QUA, Qi-EP, Qi-CON, Qi-HAL)	this paper
<i>Four axioms (AX-1–AX-4)</i>		
AX	Axiom (axiomatic primitive; non-negotiable)	GG-A1-Constitution, §1.5
LC	AX-1 Locality and Causality	GG-A1-Constitution, §2.2
GGEP	AX-2 Geometry–Gauge Equivalence Principle	GG-A1-Constitution, §2.3
OGGEp	AX-3 Operational/Open Geometry–Gauge Equivalence Principle (boundary observation dictionary)	GG-A1-Constitution, §2.4
MSMF	AX-4 Maximum Symmetry, Minimum Freedom (no hidden continuous knobs)	GG-A1-Constitution, §2.5
<i>Seven postulates (POS-1–POS-7)</i>		
POS	Postulate (admissibility filter; testable in principle)	GG-A1-Constitution, §1.5
GI	POS-1 Gauge invariance across scales (LGI / DGI / GGI)	GG-A1-Constitution, §3.3
UEP	POS-2 Unified Extremal Principle (PLA $\Leftrightarrow$ PLT / MEPP / FEP)	GG-A1-Constitution, §3.4
AF	POS-3 Anomaly Freedom and cancellability by inflow	GG-A1-Constitution, §3.5
CP	POS-4 Closure and Positivity (CPTP-admissible open dynamics)	GG-A1-Constitution, §3.6
RG	POS-5 Renormalization-Group covariance (scale covariance)	GG-A1-Constitution, §3.7
OGN	POS-6 Occam Gauge Normalization (canonical integer/rational normalization)	GG-A1-Constitution, §3.8

continued on next page

Table 55 continued from previous page

Symbol	Meaning	Primary link
MF	POS-7 Measurement and Falsifiability (operational contract + kill criteria)	GG-A1-Constitution, §3.9
LGI	Local Gauge Invariance (microscopic / closed-system layer of GI)	GG-A1-Constitution, §3.3
GGI	Global Gauge Invariance (macroscopic / ledger layer of GI)	GG-A1-Constitution, §3.3
<i>Gi-4 pillars</i>		
TGG	Total Geometric–Gauge object (Pillar I of Gi-4)	GG-A6-GRIP
CCI	Covariant Conservation Identity (Pillar II of Gi-4)	GG-A6-GRIP
GFF	Geometric Free-Energy Functional (Pillar III of Gi-4)	GG-A6-GRIP
GRIP	Gauge-Runtime Invariance Pipeline (Pillar IV of Gi-4)	GG-A6-GRIP
GEP	Geometry–Gauge Energy Principle (energetic block of GFF)	GG-A6-GRIP
CEP	Curvature–Entropy Principle (entropic block of GFF)	GG-A6-GRIP
DGI	Dynamical Gauge Invariance (initial stage of GRIP)	GG-A6-GRIP
GSSB	Geometric Spontaneous Symmetry Breaking (middle stage of GRIP)	GG-A6-GRIP
GA	Geometric Attractor (final stage of GRIP)	GG-A6-GRIP
GGA	Ground Geometric Attractor: the unique deepest GA of an admissible stratum, characterized by the four pillars in Theorem 62.1	GG-A6-GRIP
GRIP-In	Fixed-landscape regime of the runtime (descent halting at the GGA)	GG-A6-GRIP
GRIP-Out	Landscape-reorganizing regime (supplies the next descent’s initial datum; internalized by the Qi-HAL switch)	GG-A6-GRIP; GG-A11-Origin
<i>Qi-4 pillars</i>		
Qi-QUA	Quantum information for QGL (Quadruple Gauge Locking) – Pillar I of Qi-4	this paper
Qi-EP	Quantum information for the UEP (Unified Extremal Principle) – Pillar II of Qi-4	this paper
Qi-CON	Quantum information for Conformal Invariance – Pillar III of Qi-4	this paper
Qi-HAL	Quantum information for the Holographic Attractor Lattice – Pillar IV of Qi-4	this paper
OH	Outer Hologram (boundary projection of bulk data; outward face of EIC)	this paper
IH	Inner Hologram (inner self-modeling layer; inward face of EIC)	this paper
EIC	Evolving Inner Cosmos (boundary completion of Qi-4 as the OH+IH composite)	this paper
<i>Qi-HAL acronyms</i>		
MePMoS	Memory–Prediction–Motion Sensing	this paper, Qi-HAL
NHT	Neural Holographic Tomography	this paper, Qi-HAL

continued on next page

Table 55 continued from previous page

Symbol	Meaning	Primary link
HAL	Holographic Attractor Lattice	this paper, Qi-HAL
BFFT	Binary Fourier / Fractal Fourier Transform	this paper, Qi-HAL
<i>GDOS-internal acronyms (program-defined)</i>		
SIP	System Instance Packet (canonical GDOS input unit)	this paper, TSIP (T-Qi4-8.2)
SeOS	Selfish Open System	GG-A5-GDOS + this paper
<i>Program-defined mathematical symbols</i>		
Q_B	Quantum-information Bit – the energy–information Arisaka constant $Q_B \approx 0.1555$ eV per qubit; one ATP carries two, so $2N_A Q_B$ matches the standard ATP-hydrolysis free energy at biological temperatures	this paper, Definition 14.2; bulk origin GG-A2-GG6 (T8-GG6)
<i>Standard physics abbreviations (plain Roman, journal convention)</i>		
<i>Open-system / quantum-channel formalism</i>		
CPTP	Completely Positive Trace-Preserving (channel admissibility for open evolution; AX-3/POS-4)	—
GKLS	Gorini–Kossakowski–Lindblad–Sudarshan (master equation for Markovian CPTP semigroups; POS-4)	—
GNS	Gelfand–Naimark–Segal (construction)	—
NZ	Nakajima–Zwanzig (memory-kernel projection for non-Markovian open dynamics)	—
POVM	Positive Operator-Valued Measure (generalized measurement model; AX-3/POS-7)	—
<i>Quantum and classical foundational theories</i>		
BRST	Becchi–Rouet–Stora–Tyutin (cohomology)	—
GR	General Relativity	—
QCD	Quantum Chromodynamics	—
QED	Quantum Electrodynamics	—
QFT	Quantum Field Theory	—
QM	Quantum Mechanics	—
<i>Standard Model and beyond</i>		
EFT	Effective Field Theory	—
EW	Electroweak	—
GUT	Grand Unified Theory	—
SM	Standard Model of particle physics	—
SSB	Spontaneous Symmetry Breaking (closed-system; generalized by GSSB)	—
SUSY	Supersymmetry	—

continued on next page

Table 55 continued from previous page

Symbol	Meaning	Primary link
<i>Higher-dimensional constructions</i>		
AdS	Anti-de Sitter spacetime	—
AdS/CFT	Anti-de Sitter / Conformal Field Theory correspondence	—
CFT	Conformal Field Theory	—
KK	Kaluza–Klein (mode expansion / dimensional reduction)	—
RS	Randall–Sundrum (warped extra-dimension model)	—
<i>Quantum gravity and cosmology</i>		
Λ CDM	Standard cosmological model with cosmological constant + cold dark matter	—
LQG	Loop Quantum Gravity	—
TOV	Tolman–Oppenheimer–Volkoff (equation)	—
<i>Foundations of quantum mechanics</i>		
CHSH	Clauser–Horne–Shimony–Holt (inequality)	—
EPR	Einstein–Podolsky–Rosen (paradox / argument)	—
<i>Open-system / extremal-principle sub-acronyms</i>		
FEP	Free Energy Principle (boundary inference avatar of UEP)	this paper, Qi-EP
MEPP	Maximum Entropy Production Principle (heuristic boundary avatar of UEP)	this paper, Qi-EP
PLA	Principle of Least Action (closed/bulk avatar of UEP; standard variational dynamics)	this paper, Qi-EP
PLT	Principle of Least Time (operational/boundary avatar of UEP; fastest lawful path)	this paper, Qi-EP

Figure 17: hierarchy of program-defined acronyms in the GG-Theory program. Top: the framework name (GG-Theory) and its constitutional charter (GG-Constitution). Tier 2: the four axioms (AX-1–AX-4) and seven postulates (POS-1–POS-7) that together define the admissibility predicate. Tier 3: three bulk objects forced by the constitution — GG-6 (6D bulk), BI-6 (modified Bianchi identity), OGN-5 (integer ledger). Tier 4: the boundary law GDOS that the bulk projects to via OGGEp, with input acronym SIP. Tier 5: the two four-pillar structures, Gi-4 (geometric flank, with sub-acronyms GEP/CEP blocks of GFF and DGI/GSSB/GA stages of GRIP) and Qi-4 (quantum-information flank, with Qi-EP sub-acronyms PLT/MEPP/FEP, Qi-CON sub-concepts Conformal/Lévy/Fractal/ $1/f$, Qi-HAL sub-acronyms MePMoS/NHT/HAL/BFFT, the boundary completion OH+IH=EIC, and the runtime carrier SeOS (Selfish Open System) underneath). Bottom: the energy–information Arisaka constant Q_B (derived secondary quantity), with the Arisaka equation $J_E/Q_B = J_e = J_B = J_{Qi}$ that locks the four boundary currents through one conversion factor.




Figure 17: How the program’s own names are related to each other. The chart is a hierarchy, not a list: the charter sits at the top, the geometry it forces below it, the boundary law below that, and the derived pillars at the foot, so the position of a name on the chart tells the reader what it depends on. Names at the same level are siblings introduced by the same construction. The purpose is orientation rather than definition — the meanings are in the abbreviations table — and a reader who has lost track of which name belongs to which layer can recover it here in one glance.

B Glossary of Key Symbols

Table 56: Glossary of key mathematical symbols. Symbols are listed in the order in which the construction introduces them — the bulk and its factors first, then the coordinates, the ledger integers, the metrics and the curvatures — with a one-line meaning for each. The ordering is deliberate: read straight down, the table is a compressed statement of how the geometry is built, and a reader who has lost the thread of a derivation can usually recover it here faster than by searching back through the text.

Symbol	Meaning
$M_6, M_4, I_r, S_\theta^1$	6D bulk, 4D boundary, radial interval, compact angular direction
r, θ	Radial coordinate $\in [0, K]$, angular $\in [0, 2\pi)$
K	Dual-lock integer = 35
$g_{AB}, g_{\mu\nu}$	6D bulk metric, 4D boundary metric
R_6, R_4	6D Ricci scalar, 4D Ricci scalar
$M_{\text{Pl}}, M_6, M_{\text{lock}}, v$	Reduced Planck mass, 6D fundamental mass, KK lock scale, EW vev
Λ_4	4D cosmological constant
G_N	Newton's gravitational constant
$E_{\text{Pl}}, E_{\text{GUT}}$	Planck and GUT energy scales
$\Phi(x, r, \theta)$	Generic bulk field
$\phi^{(n, \ell)}(x)$	4D KK-mode field (radial index n , angular index ℓ)
$f_n(r)$	Radial profile of n -th KK mode
$m_{n, \ell}^2$	KK mass squared = $m_n^2 + \ell^2$ in natural units
$\mathcal{H}_{\text{Bulk}}, \mathcal{H}_{\text{Bdy}}, \mathcal{H}_{\text{env}}$	Bulk, boundary, environment Hilbert spaces
$\rho_{\text{Bulk}}, \rho_{\text{Bdy}}$	Bulk, boundary density operators
$\mathcal{P}, \mathcal{Q}$	Nakajima–Zwanzig projection superoperators
$\mathcal{L}, \mathcal{L}_{\text{GDOS}}$	Liouville superoperator, GDOS generator
$H_{\text{Bulk}}, H_{\text{eff}}$	Bulk Hamiltonian, effective boundary Hamiltonian
L_α, γ_α	Jump operators and rates of GKLS
$\Phi_{t,s}$	Boundary CPTP evolution channel
$\mathcal{I}_\alpha, E_\alpha$	Instrument and POVM effect
$K(t, s), \mathcal{I}(t)$	NZ memory kernel and inhomogeneous initial-correlation term
$\mathcal{R}$	Boundary reduction map
$\mathcal{D}_{\text{OGGEP}}$	OGGEP dictionary = $\mathcal{R} \circ \text{Tr}_{\text{env}}$
$T^{\mu\nu}, T_{\text{TGG}}^{\mu\nu}$	Boundary stress tensor, unified TGG stress tensor
$\Delta T_{\mu\nu}^{\text{GS}}, \Delta T_{\mu\nu}^{\text{KK}}$	GS-tensor and KK-tower inflow
X_4, H_3	BI-6 anomaly polynomial and Green-Schwarz two-form
$\mathcal{A}_{\text{tot}}$	Composite bulk connection
G_{tot}	Bulk gauge-gravity group $\text{Spin}(1, 5) \times \text{SU}(5) \times \text{U}(1)_B \times \text{U}(1)_{Q_i}$
Q^e, T^a, Y, B, Q_{Q_i}	SU(5) generators, hypercharge, baryon, informational charges
$F_{\mu\nu}$	Gauge field strength
s_6, s_4	6D and 4D BRST operators
$Z_4[J], W_4[J], \Gamma_4$	Generating functional, connected generating functional, effective action
$\text{Tr}_{\text{env}}, \text{Tr}_{\text{Bdy}}, \text{Tr}_{\text{Bulk}}$	Partial and full traces
$\boxed{\cdot}$	Boxed equation = key result
$\blacksquare$	End of proof

C Key equations

The following table collects the principal equations of the paper in logical order, with pointers to the sections where they are derived.

Table 57: The principal equations of the framework, collected in the order in which the argument establishes them and grouped by the paper that owns each. The first column numbers them, the second names them, and the third gives the compact form. The table is an index rather than a derivation: nothing is proved here, and each row points back to the place where the equation is obtained. It is included because a reader checking a later step usually wants the earlier equation in one line rather than in its original context.

No.	Name	Equation (compact form)
Foundation (Papers I–III)		
1	Bulk–boundary split	$\mathcal{H}_{\text{Bulk}} = \mathcal{H}_{\text{sys}} \otimes \mathcal{H}_{\text{env}}$
2	Reduced state	$\rho_{\partial}(t) = \text{Tr}_{\text{env}} \rho_{\text{Bulk}}(t)$
3	GG-6 bulk	$M_6 = M_4 \times I_r \times S_{\theta}^1$
4	BI-6 identity	$dH_3 = X_4$, locked coefficients $(1, -1, -2/5, -2/5)$
5	OGN-5 package	$(M, N, L; J, K) = (1, 3, 5; 7, 35)$
6	GKLS generator	$\dot{\rho}_{\partial} = -i[H_{\text{eff}}, \rho_{\partial}] + \sum_k (L_k \rho_{\partial} L_k^{\dagger} - \frac{1}{2} \{L_k^{\dagger} L_k, \rho_{\partial}\})$
Universal GDOS runtime (Part II)		
7	GDOS continuity	$\partial_t p(z, t) = -\nabla_A J^A(z, t)$
8	Current decomposition	$J^A = J_{\text{grad}}^A + J_{\text{curl}}^A + J_{\text{noise}}^A + J_{\text{gauge}}^A$
9	Runtime pipeline	$\text{DGI} \implies \text{GSSB} \implies \text{GA}$
10	DGI covariance	$J'^A(z') = \frac{\partial z'^A}{\partial z^B} J^B(z)$
11	GFF functional	$\mathcal{F}_{\text{GFF}}[p] = \int_{\mathcal{M}} p \Phi_{\text{GEP}} d\mu_g + \Theta_{\text{eff}} \int_{\mathcal{M}} p \ln p d\mu_g$
12	GFF monotonicity	$\dot{\mathcal{F}}_{\text{GFF}} \leq P_{\text{port}}$ (equality iff $\text{curl} = 0$ and $\text{noise} = 0$)
N	The OGN-5 ledger’s second entry, $N = 3$. Every ledger identity is written with it	
N_c	Color triplicity, the three colors of $\text{SU}(3)_C$ — what that entry counts	
N_g	The generation (family) count, 3. <i>Not</i> a ledger entry; fixed by GG-P2-N3	
N_f	The quark-flavor count of a beta function, $N_f = 2N_g = 6$ at high scale	

No.	Name	Equation (compact form)
13	SIP packet	$\text{SIP} = (\mathcal{M}, z^A, G_{AB}, \mathcal{A}_A, \mathcal{F}_{\text{GFF}}, \Sigma^{AB}, \mathcal{O})$
Qi-4 Pillars (Parts III–VI)		
14	Arisaka Ratio	$J'_E/J'_e = J'_e/J'_B = J'_B/J'_{Qi} = 1$ (physics natural units)
15	Arisaka Equation	$J_E/\mathbf{Q}_B = J_e = J_B = J_{Qi}$ (chemistry natural units)
16	Arisaka Constant	$\mathbf{Q}_B = (2/5) E_{\text{Ryd}}/K = (2/5) \times 13.6 \text{ eV}/35$
17	Qi-EP action	$\mathcal{S}[z] = \int dt \left[\frac{1}{2} G_{AB} \dot{z}^A \dot{z}^B + \dot{z}^A \mathcal{A}_A + V_{\text{diss}} + \Lambda_{\text{sel}} \right]$
18	PLT geodesic	$\ddot{z}^A + \Gamma_{BC}^A \dot{z}^B \dot{z}^C = 0$
19	Fractional generator	$\partial_t p = D_\alpha (-\Delta)^{\alpha/2} p, \quad 0 < \alpha \leq 2$
20	Two-knee spectrum	$S(f) \sim \begin{cases} S_0, & f \ll f_L \\ f^{-\beta}, & f_L \ll f \ll f_H \\ f^{-\gamma}, & f \gg f_H \end{cases}$
Persistence, completeness, and internality (Part IX)		
21	SeOS viability tube	$\mathbb{P}(\tau_{\partial \mathcal{V}} > t) \geq 1 - \delta$ for all t in the planning horizon
22	QCI definition	$\text{QCI}_S = w_1 q_{\text{Qi-QUA}} + w_2 q_{\text{Qi-EP}} + w_3 q_{\text{Qi-CON}} + w_4 q_{\text{Qi-HAL}}$
23	OH kernel	$P(s \mid z) = \text{Tr}[\mathcal{K}_{\text{OH}}(s) \rho(z)]$
24	IH reconstruction	$\partial_\tau \hat{\Phi} = -\Gamma \frac{\delta \mathcal{J}_{\text{EIC}}}{\delta \hat{\Phi}}$ (gradient descent on mismatch)
25	EIC convergence	$\dot{\mathcal{J}}_{\text{EIC}} \leq 0$ (Lyapunov decrease)
Synthesis (Part XI)		

26	Synthesis theorem	Papers I–III $\Rightarrow$ DGI+GSSB+GA $\Rightarrow$ GFF+SIP $\Rightarrow$ Qi-4 $\Rightarrow$ SeOS/QCI $\Rightarrow$ OH+IH=EIC
N	The OGN-5 ledger's second entry, $N = 3$. Every ledger identity is written with it	
N_c	Color triplicity, the three colors of $SU(3)_C$ — what that entry counts	
N_g	The generation (family) count, 3. <i>Not</i> a ledger entry; fixed by GG-P2-N3	
N_f	The quark-flavor count of a beta function, $N_f = 2N_g = 6$ at high scale	

D Technical derivations for the **GRIP** pipeline

D.1 Derivation of the transformed generator

This appendix records the algebra behind Eq. (7.2) in one place.

Let $\rho' = U\rho U^\dagger$, with $U = U_t$ differentiable in time. Then

$$\dot{\rho}' = \dot{U}\rho U^\dagger + U\dot{\rho}U^\dagger + U\rho\dot{U}^\dagger.$$

Insert $\dot{\rho} = \mathfrak{L}_t[\rho]$ and write $\rho = U^\dagger\rho'U$:

$$\dot{\rho}' = \dot{U}U^\dagger\rho' + U\mathfrak{L}_t[U^\dagger\rho'U]U^\dagger + \rho'U\dot{U}^\dagger.$$

Since $UU^\dagger = \text{id}$, differentiation gives $\dot{U}U^\dagger + U\dot{U}^\dagger = 0$, hence $U\dot{U}^\dagger = -\dot{U}U^\dagger$. Define $K_t = i\dot{U}U^\dagger$, so that $\dot{U}U^\dagger = -iK_t$. Then

$$\dot{\rho}' = U\mathfrak{L}_t[U^\dagger\rho'U]U^\dagger - i[K_t, \rho'].$$

Therefore

$$\mathfrak{L}_t^{(U)} = \mathcal{U}_t \circ \mathfrak{L}_t \circ \mathcal{U}_t^{-1} - i[K_t, \cdot].$$

This is the precise operator meaning of time-dependent gauge covariance on the boundary.

D.2 Quotient reduction of the continuity equation

Let the effective current be

$$J^A = J_\perp^A + R^A{}_\alpha \lambda^\alpha p,$$

with $J_\perp^A$ horizontal and $R^A{}_\alpha$ spanning the gauge orbits. For a gauge-invariant observable O ,

$$\frac{d}{dt}\langle O \rangle = \int (\nabla_A O) J^A d\mu = \int (\nabla_A O) J_\perp^A d\mu + \int (R_\alpha O) \lambda^\alpha p d\mu.$$

The second term vanishes identically. Hence the reduced physical continuity equation is obtained by pushing forward only the horizontal current:

$$\partial_t p_{\text{phys}} + \tilde{\nabla}_i \tilde{J}^i = 0.$$

This is why all instability and attractor analyses relevant to GSSB and GA must be performed on $\mathcal{X}_{\text{phys}}$, not on the unreduced state space.

D.3 Local bifurcation algebra

For the local scalar model

$$\mathcal{F}(q) = \mathcal{F}_0 + \frac{1}{2}aq^2 + \frac{1}{4}bq^4, \quad b > 0,$$

the stationary points satisfy $q(a + bq^2) = 0$. The Hessian is

$$\mathcal{F}''(q) = a + 3bq^2.$$

Hence:

$$\mathcal{F}''(0) = a, \quad \mathcal{F}''(q_{\pm}) = a + 3b\left(-\frac{a}{b}\right) = -2a.$$

Thus for $a > 0$, $q = 0$ is a local minimum and the broken branches do not exist; for $a < 0$, $q = 0$ is a local maximum and $q_{\pm}$ are minima. The stochastic stationary density is

$$p_*(q) \propto \exp\left[-\frac{2}{\sigma^2}\mathcal{F}(q)\right],$$

which changes from unimodal to bimodal across the same threshold. This is the simplest analytic model of open-system GSSB.

D.4 Catalog of possible Geometric Attractors

The GA definition intentionally allows several kinds of asymptotic objects. Table 58 records the most important cases.

Table 58: Catalog of the attractor types the definition admits. Six types run down the first column — fixed point, limit cycle, torus, localized bump, invariant measure and metastable branch — with the mathematical form and typical interpretation of each. The catalog matters because the framework’s terminus is often read as a fixed point and the definition is wider than that: a stable bump on a ring is an attractor in exactly the same sense, which is what lets the vocabulary reach a navigation circuit without being stretched.

GA type	Mathematical form	Typical interpretation
Fixed-point GA	Stable critical point of the reduced drift	Stable phase, homeostatic set-point, committed developmental state
Limit-cycle GA	Stable periodic orbit	Molecular machine cycle, circadian rhythm, locomotor pattern, economic oscillation
Torus or quasiperiodic GA	Invariant torus with several frequencies	Multioscillator coordination, coupled circadian or sensorimotor rhythms
Localized bump GA	Structured invariant profile on a ring or sheet	Ring-attractor head-direction state, working-memory packet, spatial focus of activity
Invariant-measure GA	Stationary probability measure with structured support	Nonequilibrium steady state, critical or near-critical open system
Metastable GA	Long-lived quasi-invariant set	Decision state, learned policy, memory trace, developmental canalization

These categories are not separate theories. They are alternative concrete realizations of the same general definition once the DGI and GSSB filters have been applied.

E Technical derivations for Qi-QUA

E.1 Coefficient-space to natural-unit conversion

This appendix records the sector-rescaling algebra behind Theorem 15.3.

Write the magnitude coefficient-space vector as

$$\mathbf{v}' = \begin{pmatrix} |J'_E| \\ |J'_e| \\ |J'_B| \\ |J'_{Qi}| \end{pmatrix} = r \begin{pmatrix} 1 \\ 1 \\ 2/5 \\ 2/5 \end{pmatrix}. \quad (\text{E.1})$$

Define the natural-unit conversion map

$$U_{\text{Qi-QUA}} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 5/2 & 0 \\ 0 & 0 & 0 & 5/2 \end{pmatrix}, \quad (\text{E.2})$$

acting on the event-count vector $(J_E/\mathbf{Q}_B, J_e, J_B, J_Q)^\top$. Applying $U_{\text{Qi-QUA}}$ to Eq. (E.1) gives

$$U_{\text{Qi-QUA}} \mathbf{v}' = r \begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \end{pmatrix}. \quad (\text{E.3})$$

Equation (E.3) is exactly the event-count form of the Arisaka Equation. Multiplying the first component by $\mathbf{Q}_B$ restores the physical energy current J_E .

The same result can be read backwards. Given a natural-unit event count r , the corresponding coefficient-space magnitudes are

$$|J'_E| = r, \quad |J'_e| = r, \quad |J'_B| = \frac{2}{5}r, \quad |J'_{Qi}| = \frac{2}{5}r. \quad (\text{E.4})$$

Hence the coefficient-space and natural-unit presentations are strictly equivalent once the unit conventions are fixed.

E.2 Numerical ledger of the Arisaka Constant

For convenience we collect the basic numerical conversions of $\mathbf{Q}_B$ in one place.

Starting from

$$E_{\text{Ryd}} \approx 13.6057 \text{ eV}, \quad K = 35,$$

one finds

$$\delta E_{\text{geom}} = \frac{E_{\text{Ryd}}}{K} \approx 0.3887 \text{ eV}.$$

Applying the BI-6 magnitude ratio $2/5$ gives

$$Q_B = \frac{2}{5} \delta E_{\text{geom}} \approx 0.1555 \text{ eV}.$$

Converting to chemical units using $1 \text{ eV} = 96.485 \text{ kJ mol}^{-1}$ yields

$$Q_B \approx 15.0 \text{ kJ mol}^{-1}.$$

At $T = 310 \text{ K}$, where $k_B T \approx 0.0267 \text{ eV}$, one obtains

$$Q_B \approx 5.82 k_B T.$$

Finally, since $k_B T \ln 2 \approx 0.0185 \text{ eV}$ at 310 K ,

$$Q_B \approx 8.40 k_B T \ln 2.$$

Table 59: Compact conversion table for the Arisaka constant. Each row gives one quantity, its value and a comment. The table is a working reference rather than an argument: it collects the conversions between the geometric definition, the electronvolt-per-qubit form and the molar form that the chemical comparisons need, so that a reader checking the ATP anchor against a bench measurement does not have to reconstruct them.

Quantity	Value	Comment
$\delta E_{\text{geom}} = E_{\text{Ryd}}/K$	0.3887 eV	Per-step atomic geometric scale
$Q_B = (2/5)E_{\text{Ryd}}/K$	0.1555 eV	Arisaka Constant
Q_B in molar units	15.0 kJ mol^{-1}	Chemistry-natural-unit form
$Q_B/(k_B T)$ at 310 K	5.82	Physiological thermal comparison
$Q_B/(k_B T \ln 2)$ at 310 K	8.40	Comparison with Landauer erasure bound

E.3 Coarse-grained cone generated by ideal locked events

The purpose of this appendix is to formalize the idea that realistic processes are built from sums of ideal events plus overhead.

Let the ideal event vector in natural units be

$$\mathbf{u}_\star = (Q_B, 1, 1, 1).$$

Any process consisting of N ideal events and a nonnegative overhead vector $\Delta = (\Delta_E, \Delta_e, \Delta_B, \Delta_Q)$ has total ledger

$$\mathbf{L} = N\mathbf{u}_\star + \Delta.$$

If the overhead is monotone in the sense that $\Delta_Q = 0$ and $\Delta_E, \Delta_e, \Delta_B \geq 0$, then

$$\mathbf{L} = (Q_B N + \Delta_E, N + \Delta_e, N + \Delta_B, N).$$

Hence

$$L_E - Q_B L_Q = \Delta_E \geq 0, \quad L_e - L_Q = \Delta_e \geq 0, \quad L_B - L_Q = \Delta_B \geq 0.$$

This gives the inequality form quoted in Proposition 16.1.

Geometrically, the set of ideal events spans a ray, while the set of coarse-grained processes spans a convex cone whose distinguished edge is the ideal Qi-QUA ray. The closer a real process lies to that edge, the closer it is to ideal locked efficiency.

E.4 Event-budget table

Table 60 provides a quick-reference ledger converting a few simple event counts into their ideal energy floors.

Table 60: Ideal event budgets for locked boundary events. The first column counts the events, the second and third give the ideal energy in electronvolts and in molar units, and the fourth an interpretation. The row at two quanta is the one that meets experiment — it is the ATP hydrolysis figure — and the rows above and below it are included so that a reader can see how far the neighboring integers sit from the measured value, which is the sharpest available check that the multiplicity is forced rather than fitted.

Locked events	Ideal energy	Molar energy	Interpretation
1	1Q _B	15.0 kJ mol ⁻¹	One minimal ideal boundary update
2	2Q _B	30.0 kJ mol ⁻¹	Two coupled commits or one two-bit update
3	3Q _B	45.0 kJ mol ⁻¹	Representative multi-step molecular act
8	8Q _B	120.0 kJ mol ⁻¹	One byte-equivalent update in the ideal limit
N	NQ_B	$15.0N$ kJ mol ⁻¹	General linear scaling for an ideal locked process

F Technical derivations for Qi-EP

F.1 Reduction of the current-density action to a trajectory functional

This appendix spells out in detail the reduction from the current-density functional (26.6) to the sharply localized trajectory functional (27.2). The calculation is elementary but it is worth recording explicitly because it shows exactly where the effective metric and connection enter.

Let

$$\rho(\zeta, t) = \delta(\zeta - z(t)), \quad J^a(\zeta, t) = \delta(\zeta - z(t)) \dot{z}^a(t). \quad (\text{F.1})$$

For any smooth test function $f(\zeta)$,

$$\int_{\mathcal{X}} f(\zeta) \rho(\zeta, t) d\zeta = f(z(t)), \quad \int_{\mathcal{X}} f(\zeta) J^a(\zeta, t) d\zeta = f(z(t)) \dot{z}^a(t). \quad (\text{F.2})$$

Substitute Eq. (F.1) into the quadratic part of Eq. (26.6):

$$\frac{1}{4} \int_{\mathcal{X}} (J - \rho u)^a (\rho D)_{ab}^{-1} (J - \rho u)^b d\zeta = \frac{1}{4} (\dot{z} - u(z, t))^a D_{ab}^{-1}(z, t) (\dot{z} - u(z, t))^b. \quad (\text{F.3})$$

The continuity equation is also automatically satisfied:

$$\partial_t \rho + \nabla_a J^a = -\dot{z}^a \partial_a \delta(\zeta - z(t)) + \partial_a [\delta(\zeta - z(t)) \dot{z}^a] = 0. \quad (\text{F.4})$$

The projected potential $\mathcal{V}[\rho]$ reduces to its pointwise representative $V(z, t)$. Therefore

$$\mathcal{A}_{\text{eff}}[\rho, J] \longrightarrow \mathcal{A}_{\text{trj}}[z] = \int_0^T dt \left[\frac{1}{4} (\dot{z} - u)^a D_{ab}^{-1} (\dot{z} - u)^b + V(z, t) \right]. \quad (\text{F.5})$$

Now define $g_{ab} = \frac{1}{2} D_{ab}^{-1}$ and $a_a = g_{ab} u^b$. Then

$$\begin{aligned} \frac{1}{4} (\dot{z} - u)^a D_{ab}^{-1} (\dot{z} - u)^b &= \frac{1}{2} g_{ab} \dot{z}^a \dot{z}^b - g_{ab} u^a \dot{z}^b + \frac{1}{2} g_{ab} u^a u^b \\ &= \frac{1}{2} g_{ab} \dot{z}^a \dot{z}^b - a_a \dot{z}^a + \frac{1}{2} g_{ab} u^a u^b, \end{aligned} \quad (\text{F.6})$$

which is Eq. (27.5). The Euler–Lagrange equations follow by direct variation:

$$\delta \mathcal{A}_{\text{trj}} = \int_0^T dt \left[g_{ab} \dot{z}^b \delta \dot{z}^a + \frac{1}{2} (\partial_c g_{ab}) \dot{z}^a \dot{z}^b \delta z^c - (\partial_c a_a) \dot{z}^a \delta z^c - a_a \delta \dot{z}^a + (\partial_c U_{\text{eff}}) \delta z^c \right]. \quad (\text{F.7})$$

Integrating by parts the terms proportional to $\delta \dot{z}^a$ gives

$$\delta \mathcal{A}_{\text{trj}} = \int_0^T dt \left[-\frac{d}{dt} (g_{ab} \dot{z}^b - a_a) + \frac{1}{2} \partial_a g_{bc} \dot{z}^b \dot{z}^c - \partial_a a_b \dot{z}^b + \partial_a U_{\text{eff}} \right] \delta z^a, \quad (\text{F.8})$$

and setting $\delta \mathcal{A}_{\text{trj}} = 0$ yields Eq. (27.7).

The Jacobi reduction is equally standard. On a fixed energy shell

$$E = \frac{1}{2} g_{ab} \dot{z}^a \dot{z}^b + U_{\text{eff}}, \quad (\text{F.9})$$

solve for the kinetic term and write the action along the curve γ rather than as an integral over an arbitrary parameter. One then obtains Eq. (27.9). The result justifies the interpretation of PLT as the trajectory gauge of the projected 6D action.

F.2 Force-controlled extremum and the dual minimum-dissipation form

This appendix expands the current-gauge calculation and makes the relation between MEPP and the dual minimum-dissipation formulation completely explicit.

Start from the Rayleighian

$$\mathcal{R}_\rho[J; X] = \int_{\mathcal{X}} \left(J^a X_a - \frac{1}{2} J^a (L^{-1})_{ab} J^b \right) dz, \quad (\text{F.10})$$

with continuity and integrated ledger constraints imposed by Lagrange multipliers. Including a multiplier $\lambda(z, t)$ for continuity gives

$$\tilde{\mathcal{R}} = \mathcal{R}_\rho[J; X] + \int_{\mathcal{X}} \lambda (\partial_t \rho + \nabla_a J^a) dz. \quad (\text{F.11})$$

Variation with respect to J yields

$$X_a - (L^{-1})_{ab} J^b - \partial_a \lambda = 0. \quad (\text{F.12})$$

Thus the constrained stationary current is

$$J^a = L^{ab} (X_b - \partial_b \lambda). \quad (\text{F.13})$$

The multiplier λ simply enforces admissibility; it does not change the local quadratic structure of the extremum.

Evaluate the Rayleighian at the stationary point:

$$\begin{aligned} \mathcal{R}_{\rho, \star} &= \int_{\mathcal{X}} \left[J_\star^a X_a - \frac{1}{2} J_\star^a (L^{-1})_{ab} J_\star^b \right] dz \\ &= \frac{1}{2} \int_{\mathcal{X}} J_\star^a X_a dz = \frac{1}{2} \dot{\Sigma}_\star, \end{aligned} \quad (\text{F.14})$$

where the last equality uses the definition of entropy production.

Now compute the second variation. At fixed X , ρ , and constraint surface,

$$\delta^2 \mathcal{R}_\rho = - \int_{\mathcal{X}} \delta J^a (L^{-1})_{ab} \delta J^b dz \leq 0. \quad (\text{F.15})$$

Therefore the stationary point is a maximum. This is the precise force-controlled MEPP statement of the present framework.

To obtain the dual picture, Legendre-transform Eq. (F.10) with respect to the force variable. The dual functional is

$$\mathcal{D}_\rho[J; X] = \frac{1}{2} \int_{\mathcal{X}} J^a (L^{-1})_{ab} J^b dz - \int_{\mathcal{X}} J^a X_a dz. \quad (\text{F.16})$$

At the common stationary point,

$$\delta \mathcal{D}_\rho = 0, \quad \delta^2 \mathcal{D}_\rho = \int_{\mathcal{X}} \delta J^a (L^{-1})_{ab} \delta J^b dz \geq 0, \quad (\text{F.17})$$

so the same current minimizes $\mathcal{D}_\rho$. This is Onsager's least-dissipation form familiar from near-equilibrium irreversible thermodynamics: it varies the fluxes at fixed forces, and it returns the same stationary point.

Two distinct minimum statements live in this neighborhood, and they should not be run together. Onsager's least dissipation, just derived, is the *flux-controlled* Legendre transform: forces fixed, fluxes varied. Prigogine's minimum entropy production, derived in Appendix F.6(ii), is the *force-side* statement: the transport law fixed, some forces clamped and the rest free, and the stationary point is the steady state. Both are duals of the one stationary point, but of different variables, and it is the second — not the first — that maximum entropy production has been contrasted with for seventy years. The present framework accommodates all three readings because all three arise from the same projected quadratic form.

F.3 Pullback to the statistical manifold and the Fisher metric

This appendix derives the Fisher metric appearing in the inferential gauge. Let $q_\mu(\eta)$ be a smooth parametric family of densities indexed by coordinates μ^i . A small parameter displacement $\mu \rightarrow \mu + d\mu$ changes the density by

$$q_{\mu+d\mu}(\eta) = q_\mu(\eta) + \partial_i q_\mu(\eta) d\mu^i + \mathcal{O}(d\mu^2). \quad (\text{F.18})$$

The KL divergence between neighboring family members is

$$\begin{aligned} D_{\text{KL}}(q_{\mu+d\mu} \| q_\mu) &= \int q_{\mu+d\mu} \log \frac{q_{\mu+d\mu}}{q_\mu} d\eta \\ &= \frac{1}{2} g_{ij}^{\text{F}}(\mu) d\mu^i d\mu^j + \mathcal{O}(d\mu^3), \end{aligned} \quad (\text{F.19})$$

where

$$g_{ij}^{\text{F}}(\mu) = \int q_\mu(\eta) \partial_i \log q_\mu(\eta) \partial_j \log q_\mu(\eta) d\eta. \quad (\text{F.20})$$

Thus the Fisher metric is the quadratic form measuring distinguishability of neighboring internal models.

The metric is not installed by convention; it is the second variation of the ledger's own informational term, and the computation deserves its three lines. Write

$$h(\mu) = D_{\text{KL}}(q_\mu \| q_{\mu_0}) = \int q_\mu \log \frac{q_\mu}{q_{\mu_0}} d\eta, \quad (\text{F.21})$$

which is the CEP contribution to $\mathcal{F}_{\text{GFF}}$ of Eq. (8.1) evaluated on the family, up to the additive reference term. Differentiating once, by the product rule,

$$\frac{\partial h}{\partial \mu^i} = \int \partial_i q_\mu \log \frac{q_\mu}{q_{\mu_0}} d\eta + \int q_\mu \partial_i \log q_\mu d\eta, \quad (\text{F.22})$$

and the second integral vanishes identically in μ by the score identity — differentiate the normalization under the integral sign,

$$\int q_\mu \partial_i \log q_\mu d\eta = \int \partial_i q_\mu d\eta = \partial_i \int q_\mu d\eta = \partial_i(1) = 0. \quad (\text{F.23})$$

At $\mu = \mu_0$ the logarithm in the first integral of Eq. (F.22) is $\log 1 = 0$, so the whole first derivative vanishes at the base point — as it must for a divergence attaining its minimum, zero, there. The metric therefore comes from the second derivative. Differentiating Eq. (F.22) again (the score term contributes nothing, being identically zero),

$$\frac{\partial^2 h}{\partial \mu^j \partial \mu^i} = \int \partial_j \partial_i q_\mu \log \frac{q_\mu}{q_{\mu_0}} d\eta + \int \partial_i q_\mu \partial_j \log q_\mu d\eta, \quad (\text{F.24})$$

and at $\mu = \mu_0$ the first integral again dies on $\log 1 = 0$, while the second, using $\partial_i q = q \partial_i \log q$, is exactly the expectation of the product of scores. Hence

$$\frac{\partial^2}{\partial \mu^i \partial \mu^j} D_{\text{KL}}(q_\mu \| q_{\mu_0}) \Big|_{\mu=\mu_0} = g_{ij}^{\text{F}}(\mu_0), \quad (\text{F.25})$$

the Fisher metric of Eq. (F.20). The geometry that the inference literature installs by convention arrives here as the Hessian of the CEP term of the ledger — which is why the inferential gauge inherits the same budget as the other two, rather than carrying a metric of its own. Restricting the current-density action to the statistical slice $\rho_{\mu,s,a}(\eta) = q_\mu(\eta) \delta(s - s_t) \delta(a - a_t)$, whose time dependence enters through

$$\partial_t \rho = \frac{\partial q_\mu}{\partial \mu^i} \dot{\mu}^i \delta(s - s_t) \delta(a - a_t) + \dots, \quad (\text{F.26})$$

therefore converts the boundary kinetic term into the Fisher quadratic form. Therefore

$$\mathcal{A}_{\text{eff}}[\rho, J] \Big|_{s_{\text{inf}}} = \int_0^T dt \left[\frac{1}{2} g_{ij}^{\text{F}} \dot{\mu}^i \dot{\mu}^j + \mathcal{F}(s, a, \mu) \right]. \quad (\text{F.27})$$

The variational free energy itself can be decomposed as

$$\begin{aligned} \mathcal{F}(s, a, \mu) &= \int q_\mu(\eta) \log q_\mu(\eta) d\eta - \int q_\mu(\eta) \log p(\eta, s, a) d\eta \\ &= D_{\text{KL}}(q_\mu(\eta) \| p(\eta|s, a)) - \log p(s, a), \end{aligned} \quad (\text{F.28})$$

so the only μ -dependent part relevant to internal dynamics is the KL mismatch term. Taking the Euler–Lagrange equation of Eq. (F.27) and then the overdamped limit produces the natural-gradient flow

$$\dot{\mu}^i = -\kappa_\mu (g^{\text{F}})^{ij} \partial_j \mathcal{F}. \quad (\text{F.29})$$

The derivation is important because it shows that the geometry of inference is not inserted as an extra hypothesis. It is the pullback of the same projected open-system metric that already governed trajectories and currents in the previous gauges.

F.4 A two-state blanket model

A simple two-state example is useful for seeing the inferential gauge in explicit algebra. Let the hidden external variable be $\eta \in \{0, 1\}$, and parameterize the recognition family by a single scalar μ through

$$q_\mu(\eta = 1) = \sigma(\mu), \quad q_\mu(\eta = 0) = 1 - \sigma(\mu), \quad \sigma(\mu) = \frac{1}{1 + e^{-\mu}}. \quad (\text{F.30})$$

Then

$$\partial_\mu \log q_\mu(\eta) = \eta - \sigma(\mu), \quad (\text{F.31})$$

and the Fisher metric is

$$g_{\mu\mu}^F = \sum_{\eta=0}^1 q_\mu(\eta) (\eta - \sigma(\mu))^2 = \sigma(\mu)(1 - \sigma(\mu)). \quad (\text{F.32})$$

Suppose the projected generative model assigns likelihoods $p(s|\eta)$ and prior $p(\eta)$. The variational free energy is

$$\mathcal{F}(s, \mu) = q_\mu(1) \log \frac{q_\mu(1)}{p(\eta = 1, s)} + q_\mu(0) \log \frac{q_\mu(0)}{p(\eta = 0, s)}. \quad (\text{F.33})$$

Its derivative is

$$\partial_\mu \mathcal{F} = \sigma(\mu)(1 - \sigma(\mu)) \log \frac{q_\mu(1) p(\eta = 0, s)}{q_\mu(0) p(\eta = 1, s)}. \quad (\text{F.34})$$

The natural-gradient flow becomes

$$\dot{\mu} = -\kappa_\mu \frac{1}{\sigma(1 - \sigma)} \partial_\mu \mathcal{F} = -\kappa_\mu \log \frac{q_\mu(1) p(\eta = 0, s)}{q_\mu(0) p(\eta = 1, s)}. \quad (\text{F.35})$$

The stationary condition $\dot{\mu} = 0$ is

$$\frac{q_\mu(1)}{q_\mu(0)} = \frac{p(\eta = 1, s)}{p(\eta = 0, s)}, \quad (\text{F.36})$$

which simply says that the recognition odds match the posterior odds. This miniature example makes the logic transparent: the inferential gauge of Qi-EP selects the internal state whose recognition density best matches the projected generative density.

F.5 Dictionary of the three gauges

For convenience, Table 61 gives a compact dictionary of the three Qi-EP gauges.

Table 61: Dictionary of the three Qi-EP gauges. Each row is one gauge, with its state variable, the object it extremizes and its typical use. It is the appendix companion to Table 25: the same three rows, but arranged for reference rather than for argument, so that a reader working through a derivation in one gauge can find the corresponding object in another without re-reading the Part.

Gauge	State variable	Extremized object	Typical use
Trajectory gauge	$z(t)$	Generalized optical time $\mathcal{T}[z]$	Path selection, transport, signaling, branch tracking
Current gauge	(ρ, J)	Rayleighian $\mathcal{R}_\rho[J; X]$ or dual dissipation functional	Flux networks, reaction systems, metabolic or ecological throughput
Inferential gauge	$\mu(t)$	Fisher-kinetic action $\mathcal{A}_{\text{inf}}[\mu]$ with $\mathcal{F}$ potential	Adaptive systems, cells, nervous systems, cognition, action selection

The table is intentionally simple. Its purpose is to emphasize that the three languages differ in their chosen reduced variable, not in their constitutional origin. The origin is always the same projected current-density action.

F.6 The duality theorem: both readings, in full

This appendix proves Theorem 28.2. Throughout, $L = L^\top \succ 0$ is the mobility of Eq. (28.1), $\dot{\Sigma}(J) = J^a (L^{-1})_{ab} J^b$ is the flux-side entropy production, and $P(X) = X_a L^{ab} X_b$ is the force-side one.

(i) The flux side: forces clamped. The problem is to maximize $\dot{\Sigma}(J)$ subject to the power-balance constraint

$$g(J) \equiv J^a X_a - \dot{\Sigma}(J) = 0. \quad (\text{F.37})$$

Introduce a multiplier μ and extremize $\Lambda(J, \mu) = \dot{\Sigma}(J) + \mu g(J)$ freely. First the derivative of $\dot{\Sigma}$, index by index, the two pieces combining by the symmetry of L^{-1} :

$$\frac{\partial \dot{\Sigma}}{\partial J^c} = (L^{-1})_{cb} J^b + J^a (L^{-1})_{ac} = 2 (L^{-1})_{cb} J^b. \quad (\text{F.38})$$

Stationarity of Λ is then

$$2(L^{-1})_{cb} J^b + \mu(X_c - 2(L^{-1})_{cb} J^b) = 0 \implies 2(L^{-1})_{cb} J^b = \frac{\mu}{\mu - 1} X_c \quad (\mu \neq 1), \quad (\text{F.39})$$

so every candidate extremum is proportional to the physical flux,

$$J^b = k L^{bc} X_c, \quad k \equiv \frac{\mu}{2(\mu - 1)}. \quad (\text{F.40})$$

The constraint fixes k : on the candidate, $J^a X_a = k X_a L^{ab} X_b$ and $\dot{\Sigma}(J) = k^2 X_a L^{ab} X_b$, so Eq. (F.37) gives $k = k^2$, hence $k = 1$ (the root $k = 0$ is the trivial equilibrium point $J = 0$, excluded for nonzero X by maximality below). Therefore $J = LX$: the linear law of Eq. (28.7), recovered as the constrained maximizer rather than assumed.

For the orthogonality condition, note that $\dot{\Sigma}$ is homogeneous of degree two, so Euler's homogeneous-function theorem gives

$$J^a \frac{\partial \dot{\Sigma}}{\partial J^a} = 2 \dot{\Sigma}(J); \quad (\text{F.41})$$

contracting the ansatz $X_a = \lambda \partial \dot{\Sigma} / \partial J^a$ with J^a and using the constraint,

$$\dot{\Sigma} = J^a X_a = \lambda J^a \frac{\partial \dot{\Sigma}}{\partial J^a} = 2\lambda \dot{\Sigma} \implies \lambda = \frac{1}{2}, \quad (\text{F.42})$$

fixed by the constraint and by nothing else. That the extremum is a maximum follows by expanding along the constraint surface. Parameterize $J = LX + \delta$ and impose $g(J) = 0$ exactly:

$$(LX + \delta)^a X_a = (LX + \delta)^a (L^{-1})_{ab} (LX + \delta)^b \implies \delta^a X_a = -\delta^a (L^{-1})_{ab} \delta^b, \quad (\text{F.43})$$

after canceling XLX from both sides and collecting the terms linear and quadratic in δ . Substituting this admissibility condition back into $\dot{\Sigma}$,

$$\dot{\Sigma}(LX + \delta) = X_a L^{ab} X_b + 2\delta^a X_a + \delta^a (L^{-1})_{ab} \delta^b = \dot{\Sigma}(LX) - \delta^a (L^{-1})_{ab} \delta^b \leq \dot{\Sigma}(LX), \quad (\text{F.44})$$

with equality only at $\delta = 0$ because $L^{-1} \succ 0$. The physical flux is the strict constrained maximum, which is the Ziegler statement, and it selects the constitutive relation.

(ii) The force side: constitutive law fixed. Now hold $J = LX$ fixed as the transport law and treat the forces as the variables, with $X_1, \dots, X_m$ clamped by the boundary conditions and $X_{m+1}, \dots, X_n$ free. Differentiating the force-side entropy production,

$$\frac{\partial P}{\partial X_j} = \frac{\partial}{\partial X_j} (X_a L^{ab} X_b) = 2L_{jk} X_k = 2J_j \quad (j > m), \quad (\text{F.45})$$

so P is stationary on the affine subspace of clamped forces exactly when the fluxes conjugate to the *free* forces vanish — which is the definition of the steady state under the imposed constraints. The Hessian restricted to the free block is

$$\frac{\partial^2 P}{\partial X_j \partial X_{j'}} = 2L_{jj'} \succ 0 \quad (j, j' > m), \quad (\text{F.46})$$

positive definite by the positivity of L ; the stationary point is therefore a strict minimum. This is the minimum-entropy-production theorem of Prigogine (1967) in its exact scope: near-equilibrium, linear constitutive law, and a partition of the forces into clamped and free. It selects the state.

(iii) One stationary point. Both computations locate the same configuration, $J_\star = LX_\star$, and at it $\dot{\Sigma}(J_\star) = J_\star^a (L^{-1})_{ab} J_\star^b = X_{\star a} L^{ab} X_{\star b} = P(X_\star)$. The two extremal characterizations differ only in which variables are varied and which data are held fixed: (i) varies fluxes at fixed forces and returns the transport law; (ii) varies free forces at fixed transport law and returns the state. They are Legendre-dual readings of the single stationary point of the Rayleighian (28.5). $\square$

F.7 The failure theorem: the exact obstruction

This appendix proves Theorem 30.1. Split the lawful drift of Eq. (4.4) into its dissipative part and the two retained sectors,

$$\rho u^a = \rho u_d^a + \rho u_c^a + \rho u_g^a, \quad (\text{F.47})$$

where u_d collects the gradient and noise sectors, u_c is the divergence-free circulation, and u_g generates motion along the dynamically maintained gauge directions. Carry the split through the first variation of the transport cost of Eq. (26.6) (the constant factor converting $(\rho D)^{-1}$ to L^{-1} is absorbed throughout):

$$\delta \mathcal{A}_{\text{eff}} = \frac{1}{2} \int_{\mathcal{X}} (J - \rho u_d)^a (L^{-1})_{ab} \delta J^b dz - \frac{1}{2} \int_{\mathcal{X}} (\rho u_c + \rho u_g)^a (L^{-1})_{ab} \delta J^b dz. \quad (\text{F.48})$$

The first integral is the Rayleighian variation of Eq. (28.6); the second is the cross term that the quadratic reduction discarded. Setting the total to zero for all admissible δJ gives the physical stationary current (30.1).

Now test whether any entropy-production functional is extremal there. Expand $\dot{\Sigma}$ about $J_\star$, using Eq. (F.38):

$$\dot{\Sigma}(J_\star + \delta) - \dot{\Sigma}(J_\star) = 2 \delta^a (L^{-1})_{ab} J_\star^b + \delta^a (L^{-1})_{ab} \delta^b = 2 \delta^a X_a + 2 \delta^a (L^{-1})_{ab} (\rho u_c + \rho u_g)^b + \delta^a (L^{-1})_{ab} \delta^b. \quad (\text{F.49})$$

The first term is removed by admissibility exactly as in Eq. (F.43); the third is the benign quadratic that made the flux-side reading a maximum. The middle term is the obstruction. It is *linear* in δ , of either sign, and it survives at $J_\star$ unless $\rho u_c + \rho u_g$ is L^{-1} -orthogonal to every admissible variation. A stationary point carrying a surviving linear term is not an extremum of anything.

That the orthogonality condition is exactly the work condition is seen by taking the admissible variations generated by shifting the ledger potential, $\delta J^a = L^{ab} \partial_b \chi$. Then

$$\delta^a (L^{-1})_{ab} (\rho u_c + \rho u_g)^b = \partial_a \chi (\rho u_c + \rho u_g)^a, \quad \chi = -\Phi \implies \int_{\mathcal{X}} (\cdots) dz = -\frac{\mathcal{C}_{\text{curl}} + \mathcal{G}_{\text{gauge}}}{\Theta_{\text{eff}}}, \quad (\text{F.50})$$

with $\Phi = \delta \mathcal{F}_{\text{GFF}} / \delta \rho$ the ledger potential. The vanishing of the cross term for this family of variations is therefore precisely the condition $\mathcal{C}_{\text{curl}} = \mathcal{G}_{\text{gauge}} = 0$ of the balance identity (8.3). Where either sector does work against the ledger, no functional of the entropy-production family is stationary at the physical state, and the flux-side and force-side readings fail together. Because the inferential gauge is obtained by restricting the same functional to a statistical slice, it inherits the same obstruction. $\square$

F.8 The overdamped limit, and why it is legitimate

The passage from the second-order inferential equation (29.7) to the first-order natural-gradient flow (29.8) is a limit, and limits deserve their hypotheses. The prototype is the familiar one from mechanics,

$$m\ddot{x} + \gamma\dot{x} + V'(x) = 0, \quad \frac{|m\ddot{x}|}{|\gamma\dot{x}|} \sim \frac{m}{\gamma\tau} \longrightarrow 0 \quad \text{as } \gamma\tau \gg m, \quad (\text{F.51})$$

a particle in molasses forgetting its inertia: on histories varying over a timescale τ long compared with m/γ , the acceleration term is down by one power of the timescale ratio and the surviving balance is friction against force.

Here the role of the friction is played by the mobility that the transport term of Eq. (26.6) carries onto the statistical manifold — the same $(\rho D)^{-1}$ weight, restricted to belief coordinates, that sets the rate κ_μ at which the registered states exchange with the declared ports. On relaxational histories, slow compared with that port-coupling time, $|\ddot{\mu}| \sim |\Delta\mu|/\tau^2$ while $|\dot{\mu}| \sim |\Delta\mu|/\tau$, so both left-hand terms of Eq. (29.7) — the inertial term and the velocity-quadratic Christoffel term — are suppressed by one power of the timescale ratio relative to the friction–gradient balance. Dropping them leaves Eq. (29.8).

The hypothesis is therefore explicit and checkable: the inferential gauge in its first-order form governs systems whose belief updates are slow compared with their port exchange. Systems driven faster than their ports retain the inertial and Christoffel terms, and their trajectories are not natural-gradient descents — a boundary of the reading that the second-order equation (29.7) states exactly, and which is invisible in presentations that begin from the first-order form.

G Technical derivations for Qi-CON

G.1 Radial integration and kernel homogeneity

This appendix records the scale-homogeneity algebra behind Eqs. (38.7)–(38.14) in one place. Start from the projected kernel

$$W_\Delta(z, z') = \int_{\zeta_{\min}}^{\zeta_{\max}} \omega(\zeta) \left(\frac{\zeta}{\zeta^2 + \ell^2} \right)^{2\Delta} d\zeta, \quad \ell = d_g(z, z'). \quad (\text{G.1})$$

Assume that in the active window the weight can be written as

$$\omega(\zeta) = \zeta^{\sigma-1} w(\zeta), \quad w(\zeta) \approx w_0. \quad (\text{G.2})$$

Then

$$W_\Delta(z, z') \approx w_0 \int_{\zeta_{\min}}^{\zeta_{\max}} \frac{\zeta^{\sigma+2\Delta-1}}{(\zeta^2 + \ell^2)^{2\Delta}} d\zeta. \quad (\text{G.3})$$

Set

$$\zeta = \ell u, \quad d\zeta = \ell du, \quad (\text{G.4})$$

to obtain

$$\begin{aligned}
W_\Delta(z, z') &\approx w_0 \int_{\zeta_{\min}/\ell}^{\zeta_{\max}/\ell} \frac{(\ell u)^{\sigma+2\Delta-1}}{(\ell^2 u^2 + \ell^2)^{2\Delta}} \ell \, du \\
&= w_0 \ell^{\sigma+2\Delta-1-4\Delta+1} \int_{\zeta_{\min}/\ell}^{\zeta_{\max}/\ell} \frac{u^{\sigma+2\Delta-1}}{(1+u^2)^{2\Delta}} \, du \\
&= w_0 \ell^{\sigma-2\Delta} \int_{\zeta_{\min}/\ell}^{\zeta_{\max}/\ell} \frac{u^{\sigma+2\Delta-1}}{(1+u^2)^{2\Delta}} \, du.
\end{aligned} \tag{G.5}$$

Now impose the separation condition

$$\zeta_{\min} \ll \ell \ll \zeta_{\max}. \tag{G.6}$$

Then the lower limit tends to 0, the upper limit tends to $+\infty$, and the integral becomes the constant

$$C_{\Delta,\sigma} = \int_0^\infty \frac{u^{\sigma+2\Delta-1}}{(1+u^2)^{2\Delta}} \, du = \frac{1}{2} B\left(\Delta + \frac{\sigma}{2}, \Delta - \frac{\sigma}{2}\right), \quad |\sigma| < 2\Delta, \tag{G.7}$$

Thus

$$W_\Delta(z, z') \approx w_0 C_{\Delta,\sigma} \ell^{-(2\Delta-\sigma)}. \tag{G.8}$$

If

$$2\Delta - \sigma = d_\chi + \alpha, \tag{G.9}$$

one obtains

$$W_\Delta(z, z') \propto \frac{1}{d_g(z, z')^{d_\chi + \alpha}}, \tag{G.10}$$

which is the kernel defining the fractional jump operator in the main text.

G.2 Characteristic function and stable-law tail

This appendix records the derivation of Eqs. (40.4) and (40.6). The one-dimensional leading Qi-CON equation is

$$\partial_t p(x, t) = -D_\alpha (-\partial_x^2)^{\alpha/2} p(x, t). \tag{G.11}$$

Using the Fourier convention

$$\hat{p}(k, t) = \int_{-\infty}^\infty e^{-ikx} p(x, t) \, dx, \tag{G.12}$$

and the multiplier property

$$\mathcal{F}\left[(-\partial_x^2)^{\alpha/2} p\right] = |k|^\alpha \hat{p}(k, t), \tag{G.13}$$

Eq. (G.11) becomes

$$\partial_t \hat{p}(k, t) = -D_\alpha |k|^\alpha \hat{p}(k, t). \tag{G.14}$$

This is a first-order ordinary differential equation with solution

$$\hat{p}(k, t) = e^{-D_\alpha t |k|^\alpha} \hat{p}(k, 0). \tag{G.15}$$

For the initial condition $p(x, 0) = \delta(x)$, one has $\hat{p}(k, 0) = 1$, hence

$$\phi_{X_t}(k) = \hat{p}(k, t) = e^{-D_\alpha t |k|^\alpha}, \quad (\text{G.16})$$

which is the characteristic function of a symmetric α -stable law.

To see the tail, write the inverse transform formally as

$$p(x, t) = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{ikx - D_\alpha t |k|^\alpha} dk. \quad (\text{G.17})$$

For large $|x|$, the dominant contribution comes from small $|k|$. Expanding the exponential,

$$e^{-D_\alpha t |k|^\alpha} = 1 - D_\alpha t |k|^\alpha + \dots, \quad (\text{G.18})$$

and transforming term by term gives the leading nontrivial asymptotic

$$p(x, t) \sim \frac{D_\alpha t \Gamma(1 + \alpha) \sin(\pi\alpha/2)}{\pi |x|^{1+\alpha}}. \quad (\text{G.19})$$

Thus the heavy-tail exponent is $1 + \alpha$.

G.3 Box-counting derivation of the fractal dimension

This appendix spells out the short derivation of Eq. (41.5). By self-similarity,

$$X_t \stackrel{d}{=} t^{1/\alpha} X_1. \quad (\text{G.20})$$

Therefore the typical spatial increment over a time step δt is

$$\delta x \sim (\delta t)^{1/\alpha}. \quad (\text{G.21})$$

To resolve the path at spatial scale ε , choose $\delta x \sim \varepsilon$, which implies

$$\delta t \sim \varepsilon^\alpha. \quad (\text{G.22})$$

Over an observation interval of length T , the number of distinguishable ε -scale segments is then

$$N_T(\varepsilon) \sim \frac{T}{\delta t} \sim T \varepsilon^{-\alpha}. \quad (\text{G.23})$$

However, no subset of a $d_{\mathcal{X}}$ -dimensional ambient manifold can require more than $\varepsilon^{-d_{\mathcal{X}}}$ boxes of side ε . Hence the actual covering exponent is the smaller of α and $d_{\mathcal{X}}$:

$$N_T(\varepsilon) \sim \varepsilon^{-D_f}, \quad D_f = \min(d_{\mathcal{X}}, \alpha). \quad (\text{G.24})$$

This is the box-counting form used in the main text.

G.4 Spectral density and the two-knee approximation

Starting from

$$S_X(\omega) = 2A_0 \int_{\lambda_{\text{IR}}}^{\lambda_{\text{UV}}} \frac{\lambda^\gamma}{\lambda^2 + \omega^2} d\lambda, \quad 0 \leq \gamma < 1, \quad (\text{G.25})$$

set $\lambda = \omega u$. Then

$$S_X(\omega) = 2A_0 \omega^{\gamma-1} \int_{\lambda_{\text{IR}}/\omega}^{\lambda_{\text{UV}}/\omega} \frac{u^\gamma}{1+u^2} du. \quad (\text{G.26})$$

Middle band

When $\lambda_{\text{IR}} \ll \omega \ll \lambda_{\text{UV}}$, the integral may be extended to $(0, \infty)$:

$$\int_0^\infty \frac{u^\gamma}{1+u^2} du = \frac{\pi}{2} \sec\left(\frac{\pi\gamma}{2}\right), \quad 0 \leq \gamma < 1. \quad (\text{G.27})$$

Hence

$$S_X(\omega) \approx A_0 \pi \sec\left(\frac{\pi\gamma}{2}\right) \omega^{\gamma-1} = \frac{\tilde{A}}{\omega^\beta}, \quad \beta = 1 - \gamma. \quad (\text{G.28})$$

Infrared side

When $\omega \ll \lambda_{\text{IR}}$, approximate $\lambda^2 + \omega^2 \approx \lambda^2$:

$$S_X(\omega) \approx 2A_0 \int_{\lambda_{\text{IR}}}^{\lambda_{\text{UV}}} \lambda^{\gamma-2} d\lambda = \text{constant}. \quad (\text{G.29})$$

Ultraviolet side

When $\omega \gg \lambda_{\text{UV}}$, approximate $\lambda^2 + \omega^2 \approx \omega^2$:

$$S_X(\omega) \approx \frac{2A_0}{\omega^2} \int_{\lambda_{\text{IR}}}^{\lambda_{\text{UV}}} \lambda^\gamma d\lambda \propto \omega^{-2}. \quad (\text{G.30})$$

Combining the three asymptotics yields the two-knee form of Eq. (42.12).

G.5 Compact diagnostic dictionary

For later use, Table 62 records the compact diagnostic dictionary of Part V.

Table 62: Compact dictionary of Qi-CON quantities. Each row is one quantity, with its symbol, its scaling form and an interpretation. The exponents are the reason the table exists: the same α appears under several names in different literatures — as a fractional exponent, a soft-mode exponent, a spectral slope and a fractal dimension — and a reader coming from any one of those literatures needs the translation before the pillar’s claim can be evaluated.

Quantity	Symbol	Scaling form	Interpretation
Fractional exponent	α	$\mathcal{L}_{\text{eff}}$ $-D_\alpha(-\Delta_g)^{\alpha/2}$	$\sim$ Sets dynamical scaling $t \sim \ell^\alpha$.

Quantity	Symbol	Scaling form	Interpretation
Soft-mode exponent	γ	$A_X(\lambda)\rho(\lambda) \sim \lambda^{\gamma-1}$	Measures the low-lying modal crowding of the selected branch.
Spectral slope	β	$\beta = 1 - \gamma$	Gives the intermediate $1/f^\beta$ exponent.
Fractal dimension	D_f	$D_f = \min(d_X, \alpha)$	Effective dimension of the trajectory range inside the active window.
Infrared knee	f_{IR}	$f_{\text{IR}} = \lambda_{\text{IR}}/(2\pi)$	Largest branch or attractor timescale.
Ultraviolet knee	f_{UV}	$f_{\text{UV}} = \lambda_{\text{UV}}/(2\pi)$	Smallest retained microscopic timescale.

H Technical derivations for Qi-HAL

H.1 Uniqueness of the MePMoS causal order

The dependency graph used in Theorem 48.2 can be written explicitly. Let the macro-variables be

$$V = \{\mathcal{M}_n, \hat{y}_n, a_n, s_{n+1}, \mathcal{M}_{n+1}\}.$$

The arrows are

$$E = \{\mathcal{M}_n \rightarrow \hat{y}_n, \hat{y}_n \rightarrow a_n, a_n \rightarrow s_{n+1}, s_{n+1} \rightarrow \mathcal{M}_{n+1}, \mathcal{M}_n \rightarrow \mathcal{M}_{n+1}, \hat{y}_n \rightarrow \mathcal{M}_{n+1}\}.$$

A topological sort must place each tail before its head. The chain

$$\mathcal{M}_n \prec \hat{y}_n \prec a_n \prec s_{n+1} \prec \mathcal{M}_{n+1}$$

is therefore forced. Any violation creates acausality or collapses prediction into postdiction. This is the graph-theoretic reason why MePMoS is the unique minimal ordering.

H.2 Metric recovery in NHT

Equation (49.1) implies

$$d\phi_i = \omega_i dt + \frac{2\pi}{L_i} dy^i + d\lambda.$$

Subtracting the trivial carrier term and gauge term gives

$$d\tilde{\phi}_i = \frac{2\pi}{L_i} dy^i.$$

Hence

$$dy^i = \frac{L_i}{2\pi} d\tilde{\phi}_i, \quad ds^2 = \sum_i (dy^i)^2 = \sum_i \left(\frac{L_i}{2\pi} \right)^2 (d\tilde{\phi}_i)^2.$$

For a general latent metric $g_{ij}(y)$, choose a local frame $e^a_i(y)$ with $g_{ij} = e^a_i e^a_j$ and encode the phase coordinates in that frame. The phase metric then reconstructs the curved geometry locally.

H.3 Frequency-independent **HAL** readout

Starting from

$$W_{ij} = e^{i(\phi_i - \phi_j)}, \quad x_i^r(t) = e^{i(\omega_r t + \psi_i)},$$

one computes

$$\begin{aligned} \mathcal{O}[\psi|W] &= \sum_{i,j} \overline{x_i^r(t)} W_{ij} x_j^r(t) \\ &= \sum_{i,j} e^{-i(\omega_r t + \psi_i)} e^{i(\phi_i - \phi_j)} e^{i(\omega_r t + \psi_j)} \\ &= \sum_{i,j} e^{i[(\phi_i - \phi_j) - (\psi_i - \psi_j)]}. \end{aligned}$$

The read carrier cancels exactly. Therefore ideal recall depends only on relational phase structure. Frequency dependence enters only through nonideal hardware, finite-bandwidth envelopes, or additional amplitude dynamics.

H.4 The dyadic ladder and level count of **BFFT**

If the conformal window spans $[f_{\text{IR}}, f_{\text{UV}}]$ and neighboring levels satisfy $f_{n+1} = 2f_n$, then

$$f_n = 2^n f_0.$$

Choosing $f_0 \approx f_{\text{IR}}$, the largest admissible level index satisfies

$$2^{N_{\text{lev}}-1} f_{\text{IR}} \lesssim f_{\text{UV}},$$

so

$$N_{\text{lev}} \lesssim 1 + \log_2 \left(\frac{f_{\text{UV}}}{f_{\text{IR}}} \right).$$

Taking the integer part yields Eq. (51.7). This makes the number of architectural levels a measurable consequence of the observed two-knee spectrum.

H.5 Normal-form bifurcation for the **DGI–GSSB** switch

For the scalar version of Eq. (52.1),

$$V(m) = \frac{1}{2} a m^2 + \frac{b}{4} m^4 - h m, \quad b > 0,$$

stationary points satisfy

$$am + bm^3 - h = 0.$$

At $h = 0$, one has $m = 0$ for all a , but this point is stable only for $a > 0$. When $a < 0$, the nonzero branches

$$m_{\pm} = \pm \sqrt{-\frac{a}{b}}$$

appear and are stable. This is the usual pitchfork bifurcation. Small $h \neq 0$ unfolds the pitchfork and selects one branch. In the present architecture a is controlled by accumulated evidence, so the bifurcation becomes the mathematical normal form of exploration-to-commitment dynamics.

I Technical details for SIP and GFF

I.1 Equivalence of folded and unfolded packet forms

The folded five-field packet and the more explicit unfolded runtime data are equivalent. Starting from Eq. (9.2), define the grouping map

$$\mathfrak{G} : \widehat{\text{SIP}} \mapsto [(\mathcal{Z}, g), \mathcal{P}, (\mathcal{L}_{\text{GDOS}}, \eta), (\mathcal{F}_{\text{GFF}}, \mathcal{J}), (\mathcal{O}, \mathfrak{R})].$$

The inverse map simply unpacks the generator bundle into mobility, curl, diffusion, jump, and gauge data, and unpacks the observable bundle into exported observables and falsifiers. No information is lost because each grouped entry is a direct product of the corresponding unfolded entries.

Table 63: The grouping map from unfolded runtime data to the canonical packet. The left column lists the data as it arises and the right the canonical field it belongs to. The map is many-to-one by design: the packet is a normal form, and the point of the table is to show that the grouping loses nothing, so that a reader who has assembled runtime data in some other order can check that it lands in the same place.

Unfolded data	Grouped canonical field
$(\mathcal{Z}, g_{AB})$	State manifold
$\mathcal{P}$	Ports
$(M^{AB}, \Omega^A, D^{AB}, R^A_{\alpha} \lambda^{\alpha}, \eta_{\text{jump}})$	Generator + closure
$(\mathcal{F}_{\text{GFF}}, \mathcal{J})$	Ledger bundle
$(\mathcal{O}, \mathfrak{R})$	Observable dictionary + kill criteria

Therefore the five-field form is not a simplification by omission, but a simplification by bundling.

I.2 Detailed derivation of the GFF balance law

Starting from Eq. (8.1), differentiate with respect to time:

$$\frac{d}{dt} \mathcal{F}_{\text{GFF}} = \int \frac{\delta \mathcal{F}_{\text{GFF}}}{\delta p} \partial_t p \, d\mu_g + \int p \partial_t \Phi_{\text{GEP}} \, d\mu_g - \Theta_{\text{eff}} \int p \partial_t \ln p_{\text{ref}} \, d\mu_g. \quad (\text{I.1})$$

Using $\partial_t p = -\nabla_A J^A$ and integrating by parts gives

$$\begin{aligned} \frac{d}{dt} \mathcal{F}_{\text{GFF}} &= \int \nabla_A \left(\frac{\delta \mathcal{F}_{\text{GFF}}}{\delta p} \right) J^A d\mu_g - \oint_{\partial \mathcal{Z}} \frac{\delta \mathcal{F}_{\text{GFF}}}{\delta p} J^A n_A d\Sigma \\ &\quad + \int p \partial_t \Phi_{\text{GEP}} d\mu_g - \Theta_{\text{eff}} \int p \partial_t \ln p_{\text{ref}} d\mu_g. \end{aligned} \quad (\text{I.2})$$

Substituting the current split (8.2) isolates the dissipative quadratic form, the curl contribution, the gauge contribution, and the explicit port/boundary terms. The last relative-entropy reference term is absorbed into the port power because changes in p_{ref} arise from boundary condition or control updates. This yields Eqs. (8.3)–(8.4).

If the local Einstein relation $D^{AB} = \Theta_{\text{eff}} M^{AB}$ holds, then the noise contribution combines naturally with the entropic term so that Θ_{eff} acquires the interpretation of an effective information temperature. This is the most direct bridge between GFF and standard stochastic thermodynamics (Seifert, 2012).

I.3 Event-level reduction on the locked current ray

Let the current vector be written in the ordered basis (J_e, J_B, J_{Q_i}, J_E) . Under the Arisaka Equation the allowed event increments span the one-dimensional ray

$$\delta \mathcal{J} = \delta j v, \quad v = (1, 1, 1, \mathbf{Q}_B)^\top.$$

Any polynomial event cost therefore reduces to a polynomial in the single amplitude δj . For example:

$$\begin{aligned} L_1(\delta \mathcal{J}) &= a_\alpha \delta J^\alpha = (a \cdot v) \delta j, \\ L_2(\delta \mathcal{J}) &= Q_{\alpha\beta} \delta J^\alpha \delta J^\beta = (v^\top Q v) (\delta j)^2, \\ L_3(\delta \mathcal{J}) &= T_{\alpha\beta\gamma} \delta J^\alpha \delta J^\beta \delta J^\gamma = T(v, v, v) (\delta j)^3. \end{aligned}$$

Thus all admissible local event penalties are automatically scalarized once anomaly-derived locking is imposed. This is why the scalar GFF ledger and the vector current ledger are compatible rather than competing descriptions.

I.4 Proof sketch for the viability-tube criterion

The persistence theorem rests on a familiar Lyapunov argument. Define $W(t) = \mathcal{V}(\Xi_t) - \beta \int_0^t \Pi_{\text{maint}}(s) ds$. Then Eq. (58.2) implies

$$\dot{W} \leq -\kappa \mathcal{V} + c.$$

If $\Pi_{\text{maint}} > c/\beta$ on average, then W cannot grow indefinitely, and neither can $\mathcal{V}$. Hence packet distance, boundary damage, and memory deficit all remain bounded. The tube is therefore forward invariant up to the chosen tolerance. If $\Pi_{\text{maint}} < 0$ for sufficiently long times, the opposite occurs: the accumulated negative maintenance term drives W upward, and one or more packet coordinates

leave tolerance. The proof is not sensitive to whether the target attractor is fixed, periodic, or quasiperiodic; only boundedness of the chosen packet norm matters.

I.5 Compact packet cards across representative domains

Table 64: Compact cards across representative domains. Each row is one domain, from a chemical dissipative structure to a neural circuit, with its packet core, the dominant terms in its free-energy ledger, its maintenance source and its persistence status. The cards are working templates rather than results: each is a filled-in instance of the recipe of Table 49, offered so that a reader in a new domain has a worked example of the same shape to copy.

Domain	Packet core	Dominant terms	GFF	Maintenance source	SeOS status
Chemical dissipative structure	Pattern amplitude, thermal gradient, vessel geometry	Gradient entropic ening, forcing	energy, broad-boundary	External heating / cooling	Usually proto-SeOS rather than full SeOS
Protocell	Membrane state, internal metabolic packet, permeability	Assembly mixing electrochemical coupling	cost, entropy, put	Nutrient and amphiphile through-put	Full SeOS if repair margin is positive
Bacterium	Membrane, metabolism, adaptive memory, state	Nutrient ATP ledger, formation update cost	uptake, in-	Continuous nutrient uptake and repair metabolism	Full SeOS
Neural circuit	Recurrent synaptic memory, phase code	Coherence plasticity metabolic supply	cost, cost,	Organism-level metabolic sensory ports	Full SeOS; often Qi-HAL-rich
Institution	Archive, policy state, communication network, resource stock	Coordination cost, archive maintenance, informational entropy		Resource inflow and repair of structure/records	Distributed higher-order SeOS

J Technical derivations for OH, IH, and EIC

J.1 Derivation of the outer-hologram kernel from instruments

This appendix records the instrument-theoretic derivation behind Eqs. (64.1)–(64.3).

Let the effective reduced bulk-to-boundary map be represented by a completely positive trace-preserving map $\Phi_{t,0}$, and let $\{\mathcal{I}_y\}$ be an admissible instrument family satisfying

$$\mathcal{I}_y \geq 0, \quad \int_{\mathcal{Y}} \mathcal{I}_y \, dy = I_{\partial}.$$

For a pure bulk representative $|X\rangle$, define

$$K_{\text{OH}}(y|X; t) = \text{Tr}[\mathcal{I}_y \Phi_{t,0}(|X\rangle\langle X|)].$$

Positivity follows because the trace of the product of two positive operators is nonnegative. Normalization follows from

$$\int_{\mathcal{Y}} K_{\text{OH}}(y|X; t) \, dy = \text{Tr}\left[\left(\int_{\mathcal{Y}} \mathcal{I}_y \, dy\right) \Phi_{t,0}(|X\rangle\langle X|)\right] = \text{Tr}[\Phi_{t,0}(|X\rangle\langle X|)] = 1,$$

because $\Phi_{t,0}$ is trace preserving. For a mixed bulk state $\rho_6(X, t)$, linearity immediately yields Eq. (64.2).

The derivation shows that the OH kernel is not a new postulate. It is the instrument-level unpacking of the OGGEp boundary law.

J.2 Observability and action-conditioned reconstruction

This appendix gives the standard algebra behind Theorem 64.5. For the linearized dynamics

$$\delta x_{t+1} = A_t \delta x_t + B_t \delta a_t, \quad \delta s_t = C_t \delta x_t,$$

one has

$$\begin{aligned} \delta s_0 &= C_0 \delta x_0, \\ \delta s_1 &= C_1 A_0 \delta x_0 + C_1 B_0 \delta a_0, \\ \delta s_2 &= C_2 A_1 A_0 \delta x_0 + C_2 A_1 B_0 \delta a_0 + C_2 B_1 \delta a_1, \end{aligned}$$

and so on. Collecting terms gives

$$\mathbf{s}_{0:T} = \mathcal{O}_T \delta x_0 + \mathbf{u}_{0:T},$$

where $\mathbf{u}_{0:T}$ contains the known action-dependent contributions. If $\mathcal{O}_T$ has full column rank, the least-squares estimate

$$\widehat{\delta x_0} = (\mathcal{O}_T^\top \mathcal{O}_T)^{-1} \mathcal{O}_T^\top (\mathbf{s}_{0:T} - \mathbf{u}_{0:T})$$

is unique. If not, the nullspace is nontrivial and the hidden state cannot be uniquely reconstructed from the available history.

The mathematical moral is simple: action changes what can be seen, and memory stores the history needed to exploit that fact.

J.3 Two-view triangulation from lower-dimensional sensory inputs

A familiar example of IH formation is the reconstruction of a 3D point from 2D observations at multiple viewpoints. Let $X \in \mathbb{P}^3$ be a homogeneous 3D point and $\tilde{u}_i \in \mathbb{P}^2$ its homogeneous image in frame i . Then

$$\lambda_i \tilde{u}_i = P_i X, \quad P_i = K_i [R_i | t_i].$$

For two or more views, one forms the linear system

$$AX = 0,$$

where A is assembled from the cross-product equations $\tilde{u}_i \times (P_i X) = 0$. Nonzero camera baseline or self-motion is required for unique depth recovery up to the usual projective gauge. This is the most elementary concrete instance of the general claim in the main text: a lower-dimensional instantaneous observation does not determine the world; the world becomes reconstructable only after motion generates multiple views.

The same algebra appears, abstractly, in chemical gradient sensing, head-direction stabilization, and other active-sensing settings. The substrate changes; the observability logic does not.

J.4 Internal measurement operator and selected content

Suppose the IH state is expanded in an orthonormal content basis $\{|c\rangle\}$:

$$\rho_{\text{IH}} = \sum_{c,c'} \rho_{cc'} |c\rangle\langle c'|.$$

An internal coarse-graining may be represented by block projectors $\Pi_\alpha = \sum_{c \in B_\alpha} |c\rangle\langle c|$, where B_α is a content block. Then

$$p(\alpha) = \text{Tr}(\Pi_\alpha \rho_{\text{IH}}) = \sum_{c \in B_\alpha} \rho_{cc}.$$

If the DGI→GSSB→GA sequence suppresses all but one dominant block B_{α_*} , the internal post-selected state becomes

$$\rho_{\text{IH}}^{(\alpha_*)} = \frac{\Pi_{\alpha_*} \rho_{\text{IH}} \Pi_{\alpha_*}}{\text{Tr}(\Pi_{\alpha_*} \rho_{\text{IH}})}.$$

A subsequent write or action-bias map $\mathcal{W}$ produces

$$m_{t+1} = \mathcal{W}[\rho_{\text{IH}}^{(\alpha_*)}],$$

which makes the selected content causally effective in the next MePMoS cycle. This is the algebraic heart of the internal-measurement idea.

J.5 Lyapunov decrease of the EIC functional with bounded fluctuations

For the noisy gradient flow

$$\dot{\theta} = -G_\theta^{-1} \nabla_\theta \mathcal{F}_{\text{EIC}} + \eta_t,$$

one has

$$\dot{\mathcal{F}}_{\text{EIC}} = -\nabla_{\theta} \mathcal{F}_{\text{EIC}}^{\top} G_{\theta}^{-1} \nabla_{\theta} \mathcal{F}_{\text{EIC}} + \nabla_{\theta} \mathcal{F}_{\text{EIC}}^{\top} \eta_t.$$

If η_t is zero mean with covariance bounded by Σ_{η} , the expected drift satisfies

$$\mathbb{E} \dot{\mathcal{F}}_{\text{EIC}} \leq -\lambda_{\min}(G_{\theta}^{-1}) \mathbb{E} \|\nabla_{\theta} \mathcal{F}_{\text{EIC}}\|^2 + \mathbb{E} \|\nabla_{\theta} \mathcal{F}_{\text{EIC}}\| \|\eta_t\|.$$

Whenever the deterministic descent term dominates the fluctuation term, the expected EIC functional decreases. In practice this is the appropriate statement for open systems subject to bounded noise: exact monotonicity may be broken locally, but average descent toward a basin is preserved.

J.6 Representative OH/IH/EIC cards

Table 65: Compact cards across representative systems. Each row is one system, from a molecular rotor to a rodent navigation system, with what plays the part of the outer hologram, what plays the inner one, and the resulting reading. The cards are ordered by increasing memory depth, and the value of the ordering is that it shows the inner column filling in gradually rather than switching on.

System	OH	IH	EIC reading
Molecular rotor	External torque, chemical potentials, occupancy transitions at the boundary	Conformational phase / occupancy state only	Minimal or degenerate; local cycle knowledge rather than a genuine world model
Chemotactic bacterium	Sampled ligand concentration along the trajectory	Adaptive methylation state encoding recent gradient history	Proto-EIC: local chemical world reconstructed enough to bias future motion
Ring-attractor navigation system	Visual, vestibular, and proprioceptive heading evidence	Persistent heading bump and related attractor manifold	Allocentric heading EIC with repeated correction by OH events
Rodent navigation system	Multimodal environmental stream with self-motion updates	Place/head-direction/time-structured spatial model	Rich navigational EIC; imagination and replay become possible
Human scientific cognition	Experimental data, language, symbolic and sensory streams	Multiscale internal model spanning space, time, causality, and abstraction	Consciousness-grade EIC; the system measures its own model and can report it symbolically

K Pagoda hierarchy and cross-domain technical details

K.1 Closure defects, effective floors, and error budgets

This appendix expands Definition 66.1. Let $z^{(\ell+1)} = C_{\ell+1}(z^{(\ell)})$ and define the exact induced derivative and closure defect by Eq. (66.1). A convenient normalized defect index is

$$\varepsilon_{\text{cl}}^{(\ell+1)} := \frac{\langle \|\Delta_{\ell+1}\|^2 \rangle^{1/2}}{\langle \|F_{\text{eff}}^{(\ell+1)}\|^2 \rangle^{1/2} + \delta}, \quad \delta > 0. \quad (\text{K.1})$$

The index is small when the floor is dynamically autonomous. In data-driven applications one can estimate $F_{\text{eff}}^{(\ell+1)}$ from reduced trajectories and treat ε_{cl} as a model-selection criterion for candidate higher-level packets.

A second useful quantity is the *port novelty index*

$$\nu_{\text{port}}^{(\ell+1)} := \dim \text{span } \mathcal{P}^{(\ell+1)} - \dim \text{span } \mathcal{P}_{\text{pull}}^{(\ell \rightarrow \ell+1)}. \quad (\text{K.2})$$

If $\nu_{\text{port}}^{(\ell+1)} = 0$, then the higher floor introduces no genuinely new mode of exchange and is usually better interpreted as a descriptive convenience rather than a new scientific level.

K.2 Hitting-time proof of conditional inevitability

Theorem 66.4 may be rewritten using visit counts. Let $N_W(T)$ be the number of visits to a viable window W by time T , and let S_n be the indicator that the n -th visit stabilizes. Under the assumptions of the theorem,

$$\mathbb{P}(S_n = 1 \mid \text{visit}_n) \geq p > 0.$$

Hence

$$\mathbb{P}(\text{no stabilization up to visit } N) \leq (1 - p)^N.$$

If $N_W(T) \rightarrow \infty$ almost surely as $T \rightarrow \infty$, then

$$\mathbb{P}(\tau_W = \infty) = \lim_{N \rightarrow \infty} (1 - p)^N = 0.$$

The theorem therefore uses only three ingredients: positive-measure windows, recurrent exploration, and positive stabilization probability.

K.3 Compact domain cards and observables

Table 66: Compact application cards for representative domains. Each row is one domain, with its key packet variables, candidate observables, the primary memory variable and the single most diagnostic test. The last column is the one to read first: it names, for each domain, the measurement whose outcome would most quickly distinguish a system running this runtime from one that merely looks as though it does.

Domain	Key packet variables	Candidate observables	Primary memory variable	Most diagnostic test
Catalytic chemistry	concentrations, occupancy, rotor angle	product counts, spectra, time, fluxes	conformational state	event-ledger closure plus burst statistics
Chemotactic cell	receptor state, internal methylation, trajectory	run/tumble statistics, local gradient, adaptation curve	adaptive receptor memory	motion-conditioned predictive quality
Organismal physiology	nutrient damage, flux, internal model	store, repair traces, markers	metabolic rate, sensory repair	regulatory long-term logical memory
Navigation circuit	neural manifold, latent heading/-place state, action output	spikes/LFP, movement, sensory cues, reports	recurrent attractor state	OH/IH model-improvement under controlled measurement practice
Scientific institution	agent states, archive, instruments, communication graph	data releases, publications, consensus updates	distributed archive and trained priors	SeOS positive inequality

K.4 Minimal dynamical forms across domains

For convenience, the minimal application equations from the main text are collected here in one place.

$$\begin{aligned}
\text{Chemistry: } \dot{n}_i &= \sum_r \nu_{ir} R_r(n, m) - \nabla \cdot J_i + I_i - O_i + \xi_i, & \dot{m} &= f_m(n, m) + \eta_m, \\
\text{Life: } \dot{N} &= J_{\text{in}} - J_{\text{use}} - J_{\text{leak}}, & \dot{B} &= Y J_{\text{use}} - \delta B - c_r(D), \\
&\dot{D} = \sigma - r, & \dot{m} &= f_m(y, a, m), \\
\text{Brain: } \tau_x \dot{x} &= -x + W\phi(x) + Uy + Va + \zeta_x, & \tau_m \dot{m} &= f_m - K \nabla_m \mathcal{E}_{\text{EIC}}, \\
\text{Society: } \dot{x}_i &= f_i(x_i, A, d_i) + \sum_j W_{ij} \Psi(x_j - x_i) + \eta_i, & \dot{A} &= \Gamma(\{x_i\}, d) - \Lambda A.
\end{aligned}$$

These forms are not intended as finished empirical models. They are canonical placeholders showing how the same runtime grammar appears in very different domains.

L Falsifiability summary

L.1 Compact falsifiability matrix

Table 67: Compact falsifiability matrix for the integrated framework. Each row is one runtime block, with an observable proxy for it, a representative domain in which that proxy is measurable, and the failure mode that would refute the block. Read beside Table 54 it is the same program arranged by structure rather than by signature, so that a reader who wants to know what a given pillar risks can find it in one row.

Runtime block	Observable proxy	Representative domain	Failure mode
Qi-QUA	event-level energy/information accounting	chemistry, molecular motors	no coherent locked ledger
Qi-EP	path-selection under perturbation	reaction routing, behavior, inference tasks	no common selector family
Qi-CON	spectra/tails/fractal dimension	neural activity, ecological time series, market-like records	no stable multiscale signature
Qi-HAL	memory depth and multiscale coordination	cells, brains, institutions	no measurable internal memory architecture
OH/IH/EIC	model quality under action restriction	navigation, manipulation, scientific practice	internal model unaffected by motion constraints

M Chemistry, biochemistry, and protocellular open systems

The main text argued that chemistry is the first decisive testbed for the universal-runtime claim, because chemistry is the domain in which one simultaneously sees energetic throughput, charge transfer, material reorganization, structural persistence, and information-bearing selectivity. A detector or a scattering experiment already displays lawful open-system projection, but chemistry adds a stronger demand: the system must often preserve a far-from-equilibrium structure while repeatedly performing useful cycles. That is exactly the setting in which SIP, GFF, SeOS, GRIP-GA, and the Qi-4 pillars should either become precise and useful or else fail. The point of the present appendix is therefore not decorative. It is to show, in slower detail than the main text allowed, how the GDOS runtime is expected to be instantiated in reaction networks, enzyme cycles, membrane systems, autocatalytic sets, and protocell-scale organizations (Kauffman, 1993, Landauer, 1961, Lane, 2015, Nicolis and Prigogine, 1977, Prigogine, 1967).

M.1 Why chemistry is the first nontrivial export layer

In the present program, the chemical layer is not merely “between” particle physics and biology in an informal pedagogical sense. It is the first layer at which a finite open system can both preserve a persistent identity and implement selective throughput. A molecule, membrane, or protocell must trade among at least four demands: energetic affordability, material or structural maintenance, charge or redox consistency, and information-bearing discrimination. The first two demands are familiar in thermodynamics; the third is explicit in electrochemistry; the fourth appears in selectivity, proofreading, templating, signaling, and memory-like state retention. The claim of Qi-QUA is that these demands are not independent ledgers once the underlying event is projected to the lawful boundary description of GDOS. The current-locking law should therefore be tested here first, because chemistry gives one the cleanest mesoscopic event count.

M.2 A chemistry-scale System Instance Packet

Let a chemical open system possess N_s coarse-grained species variables n_i , N_c conformational coordinates ξ_α , N_p port or membrane coordinates ϕ_a , and N_m slowly writable internal markers m_b . The reduced state coordinates are taken to be

$$z^A = (n_1, \dots, n_{N_s}, \xi_1, \dots, \xi_{N_c}, \phi_1, \dots, \phi_{N_p}, m_1, \dots, m_{N_m}). \quad (\text{M.1})$$

The corresponding runtime packet is

$$\text{SIP}_{\text{chem}} = (\mathcal{M}_{\text{chem}}, z^A, G_{AB}^{(\text{chem})}, \mathcal{A}_A^{(\text{chem})}, \mathcal{F}_{\text{GFF}}^{(\text{chem})}, \Sigma_{(\text{chem})}^{AB}, \mathcal{O}_{\text{chem}}). \quad (\text{M.2})$$

Each slot now has a direct interpretation.

First, $\mathcal{M}_{\text{chem}}$ is the accessible manifold of occupation levels, conformers, port states, and slow memory marks. It is not the full $6N$ -dimensional microscopic phase space. Second, the induced metric $G_{AB}^{(\text{chem})}$ measures the local cost or distinguishability of moving on that reduced manifold; in practice this may be inferred from susceptibility matrices, covariance matrices, or information-geometric approximations. Third, $\mathcal{A}_A^{(\text{chem})}$ is the effective gauge or curl connection that survives coarse-graining. In chemical language it packages cyclic driving, nonconservative affinities, and phase-like loop effects. Fourth, $\mathcal{F}_{\text{GFF}}^{(\text{chem})}$ is the common ledger that combines energy supply, entropy production, structural maintenance, and geometric constraints. Fifth, $\Sigma_{(\text{chem})}^{AB}$ records unresolved stochasticity from microscopic collisions, conformational substructure, and environmental variation. Finally, $\mathcal{O}_{\text{chem}}$ is the explicit dictionary of measurable outputs: concentrations, currents, rates, redox states, membrane potentials, binding fidelity, growth, division, or other experimentally accessible records.

The reduced continuity law takes the standard GDOS form

$$\partial_t p(z, t) + \nabla_A J^A(z, t) = 0, \quad J^A = J_{\text{grad}}^A + J_{\text{curl}}^A + J_{\text{noise}}^A, \quad (\text{M.3})$$

with

$$J_{\text{grad}}^A = -p M^{AB} \nabla_B \Xi_{\text{chem}}, \quad (\text{M.4})$$

$$J_{\text{curl}}^A = p \Omega^{AB} \nabla_B \Psi_{\text{chem}}, \quad (\text{M.5})$$

$$J_{\text{noise}}^A = -\frac{1}{2} \nabla_B \left(\Sigma_{(\text{chem})}^{AB} p \right). \quad (\text{M.6})$$

The first term drives relaxation on the common scalar ledger, the second encodes looped or oscillatory turnover, and the third closes positivity under unresolved fluctuations. In a simple enzyme cycle the curl term is the reduced remnant of a driven sequence of binding and release events. In a membrane-transport problem it is the remnant of circulation through a driven electrochemical loop. In an autocatalytic or templating set it captures the fact that turnover may circulate around a network rather than merely descend a scalar potential.

Table 68: One concrete interpretation of the seven packet slots in chemistry and biochemistry. The left column names the slot and the right gives its typical chemistry-scale meaning. The word *one* in the heading is deliberate: the interpretation is an instance, not a derivation, and a different chemical setting would fill the slots differently. It is included because the packet is abstract enough that a single worked filling makes the rest of the appendix legible.

SIP slot	Typical chemistry-scale meaning
$\mathcal{M}_{\text{chem}}$	Coarse-grained manifold of concentrations, conformers, port states, and slow writable markers.
z^A	Occupancies n_i , conformers ξ_α , membrane or redox variables ϕ_a , and memory-like variables m_b .
$G_{AB}^{(\text{chem})}$	Susceptibility or information metric describing the local cost of state displacement.
$\mathcal{A}_A^{(\text{chem})}$	Nonconservative affinities, cycle phases, and port couplings surviving coarse-graining.
$\mathcal{F}_{\text{GFF}}^{(\text{chem})}$	Joint ledger of free-energy supply, entropy production, structural maintenance, and geometric constraint.
$\Sigma_{(\text{chem})}^{AB}$	Stochastic forcing from collisions, conformational microstates, and environmental variation.
$\mathcal{O}_{\text{chem}}$	Measurable outputs: rates, redox or ion currents, yields, fidelity, growth, division, and robustness.

M.3 Reaction-network representation and current ledgers

Suppose the system contains reaction channels $\rho = 1, \dots, R$ with stoichiometric matrix $S_{i\rho}$ and coarse-grained fluxes j_ρ . Then

$$\partial_t n_i = \sum_{\rho=1}^R S_{i\rho} j_\rho + \eta_i, \quad (\text{M.7})$$

where η_i summarizes unresolved noise and exchange with neglected degrees of freedom. For the present program the important point is not merely stoichiometry but *ledger separation*. To each flux channel one assigns four projected event weights:

$$\Delta G_\rho, \quad q_\rho, \quad b_\rho, \quad \iota_\rho, \quad (\text{M.8})$$

interpreted respectively as energetic throughput, charged or electrochemical transfer, structural-material restructuring, and information-bearing discrimination. The associated reduced currents are then

$$J_E = \sum_{\rho} \Delta G_{\rho} j_{\rho}, \quad (\text{M.9})$$

$$J_e = \sum_{\rho} q_{\rho} j_{\rho}, \quad (\text{M.10})$$

$$J_B = \sum_{\rho} b_{\rho} j_{\rho}, \quad (\text{M.11})$$

$$J_{Qi} = \sum_{\rho} \iota_{\rho} j_{\rho}. \quad (\text{M.12})$$

At the microphysical level J_B descends from the baryonic channel of GG-6. At chemistry scale it becomes the structural-material bookkeeping current: the channel that counts net rearrangement of matter-bearing structure across the chosen boundary description. The notation is inherited, but the interpretation is mesoscopic. Likewise J_{Qi} is the effective information-throughput channel, not an arbitrary symbolic label.

A *locked elementary cycle* is defined to be a coarse-grained event class for which the four currents are proportional to a common event count n_{cyc} . In chemistry natural units one then writes

$$\frac{J_E}{Q_B} = J_e = J_B = J_{Qi} = n_{\text{cyc}}. \quad (\text{M.13})$$

Equation (M.13) is not asserted to hold for every arbitrarily coarse aggregate without correction. Rather, it is the ideal bookkeeping relation for an admissible elementary boundary update in which energy transfer, charge/redox transfer, structural reconfiguration, and information-bearing discrimination are inseparable aspects of one projected event. Deviations from exact equality are therefore diagnostic. If a biochemical model claims to implement a one-step selective event while grossly violating (M.13), then either the proposed event class is not elementary or the chosen SIP has omitted one of the relevant ports.

M.4 Qi-EP in biochemical cycles

The chemical domain is also the natural place to see why Qi-EP is not a mere restatement of one preexisting extremal principle. Consider a path $z(t)$ on $\mathcal{M}_{\text{chem}}$. The chemistry-scale GDOS path functional may be organized as

$$\mathcal{S}_{\text{chem}}[z] = \int_{t_i}^{t_f} dt \left[\frac{1}{2} g_{AB}(z) \dot{z}^A \dot{z}^B + \dot{z}^A \mathcal{A}_A^{(\text{chem})}(z) + V_{\text{diss}}(z) + \Lambda_{\text{sel}}(z, \dot{z}) \right]. \quad (\text{M.14})$$

The first term is the metric cost of motion on the reduced manifold. The second is the cycle or affinity term. The third records the dissipative penalty or throughput cost. The fourth stands for selectivity or fidelity constraints, such as proofreading, templating, or gated transport. Now consider three limiting regimes.

If the stochastic forcing is weak, the cycle structure modest, and the relevant task primarily geometric, then minimization of $\mathcal{S}_{\text{chem}}$ reduces to a shortest-time or shortest-cost path problem: this is the **PLT** face. If the task is dominated by sustained driving with strong dissipation, the extremum reorganizes into a constrained entropy-production problem: this is the **MEPP** face. If, however, the system contains internal variables m_b that encode predictions or internal templates, then part of Λ_{sel} becomes an inference-like mismatch term between internal and external port states, and the same master functional acquires a free-energy-like form: this is the **FEP** face. The physical claim of **Qi-EP** is therefore that the apparent plurality of principles emerges because different boundary tasks emphasize different sectors of the same path functional. Chemistry provides the cleanest example because the same device may operate in all three modes depending on timescale and coarse-graining.

M.5 Qi-CON signatures in chemistry and protocells

Chemistry is also where **Qi-CON** becomes sharply testable. Enzymes, ion channels, autocatalytic sets, and membrane-active protocells often display multiple relaxation scales, conformational substates, and intermittent switching. The **GDOS** expectation is that the reduced noise kernel is not generically single-scale. Instead one should often see

$$S_X(f) \sim \begin{cases} S_0, & f \ll f_L, \\ f^{-\alpha}, & f_L \ll f \ll f_H, \\ f^{-\beta}, & f \gg f_H, \end{cases} \quad (\text{M.15})$$

with f_L set by the longest architectural closure time of the system and f_H by the fastest internal port-gating or conformational switching scale. Likewise residence-time distributions may exhibit a scale window

$$P(\tau) \propto \tau^{-1-\mu}, \quad 0 < \mu < 2, \quad (\text{M.16})$$

before finite-size or finite-memory cutoffs restore integrability. In the present framework such laws are not treated as mysterious empirical curiosities. They are the expected traces of a multiscale attractor lattice whose degrees of freedom have been only partially projected to the observer's boundary.

M.6 From protocells to living systems

A protocell or minimally living system is precisely the point at which chemical runtime acquires the full **SeOS** interpretation. The boundary is no longer merely conceptual; it is a maintained membrane or other selective enclosure. The internal chemistry is no longer merely reactive; it becomes self-referential in the limited but important sense that part of the throughput is spent preserving the ports through which throughput occurs. This is the first unmistakable appearance of a selfish open system. The system persists not by violating thermodynamics but by routing gradients through a geometry that keeps its own boundary in being ([Lane, 2015](#), [Schrödinger, 1944](#)). In this sense, life at its lowest layer is a chemical proof-of-principle of the **GDOS** program.

Table 69: Chemistry-domain diagnostics. Each row is one claim of the framework, with an operational observable that bears on it and what would count as support or as tension. The third column is written to be symmetric on purpose — it says what would go against the claim as well as what would go for it — because the chemistry scale is where the framework’s one sharp number lives and it should not be tested only in the direction that confirms it.

Claim	Operational observable	What would count as support or tension
Qi-QUA current locking	Joint measurement of energy, redox/ion transfer, structural turnover, and selectivity events	Support: one elementary cycle class approximately collapses these four ledgers to a common event count. Tension: no sensible coarse-graining yields such locking.
Qi-EP unification	Fit of short-path, dissipative, and inference-like objectives within one master cycle functional	Support: regime changes can be modeled by reweighting one functional. Tension: mutually incompatible objectives are required for the same cycle class.
Qi-CON multiscale statistics	Power spectra, waiting-time tails, and state-residence histograms	Support: two knees and finite scale windows correlate with identifiable structural scales. Tension: fitted exponents float independently of architecture.
SeOS persistence	Boundary-maintenance cost vs. throughput budget	Support: a calculable maintenance sector exists and predicts loss-of-identity thresholds. Tension: persistence has no ledger representation.

M.7 Concluding remark on the chemistry layer

The strongest reason to emphasize chemistry in a landmark GDOS paper is methodological. Chemistry is the first domain in which the theory must cease to be merely a philosophical unifier and become a quantitative runtime. Concentrations, currents, binding events, redox changes, membrane potentials, yields, and fidelity measures are all observable. If the framework cannot become precise here, then the later claims about life and mind would be premature. Conversely, if it succeeds here, then the ascent to biological and cognitive systems becomes natural rather than rhetorical.

N Neural, cognitive, and consciousness-oriented open systems

The main text introduced Qi-HAL and the identity $\text{OH} + \text{IH} = \text{EIC}$ as the architectural export of GDOS to internal-model systems. The present appendix slows that argument down and states it in a more explicitly operational form. The goal is not to claim that one already possesses a final quantitative theory of consciousness. The goal is more modest and, in a way, more demanding: to show that once the GG-6/GDOS runtime is accepted, nervous systems are *not* exceptional in requiring entirely new metaphysical principles. They are exceptional only in architectural depth. The same boundary logic that governs detectors and cells now runs through large recurrent systems that sense, remember, predict, and act over many timescales (Amari and Nagaoka, 2000, Friston, 2010, Shannon, 1948, Tononi, 2004).

N.1 A neural System Instance Packet

Let a neural open system be described at reduced scale by fast activation variables x_i , phase variables θ_α , synaptic or coupling variables w_{ij} , slow memory traces m_b , and action-readiness variables a_c . One may then choose

$$z^A = (x_1, \dots, x_{N_x}, \theta_1, \dots, \theta_{N_\theta}, w_{11}, \dots, w_{N_x N_x}, m_1, \dots, m_{N_m}, a_1, \dots, a_{N_a}), \quad (\text{N.1})$$

together with the packet

$$\text{SIP}_{\text{neural}} = (\mathcal{M}_{\text{neural}}, z^A, G_{AB}^{(\text{neural})}, \mathcal{A}_A^{(\text{neural})}, \mathcal{F}_{\text{GFF}}^{(\text{neural})}, \Sigma_{(\text{neural})}^{AB}, \mathcal{O}_{\text{neural}}). \quad (\text{N.2})$$

Here $G_{AB}^{(\text{neural})}$ may be interpreted as the effective information or control metric on the reduced state space; $\mathcal{A}_A^{(\text{neural})}$ encodes phase transport, recurrent loops, and nonconservative input drive; $\mathcal{F}_{\text{GFF}}^{(\text{neural})}$ combines metabolic cost, prediction or mismatch cost, structural regularization, and task constraints; $\Sigma_{(\text{neural})}^{AB}$ records unresolved variability and fluctuating input; and $\mathcal{O}_{\text{neural}}$ specifies the readout dictionary, such as spikes, local fields, imaging signals, behavioral reports, reaction times, or action trajectories.

The universal GDOS current law again applies:

$$\partial_t p(z, t) + \nabla_A J^A(z, t) = 0, \quad J^A = -p M^{AB} \nabla_B \Xi + p \Omega^{AB} \nabla_B \Psi - \frac{1}{2} \nabla_B (\Sigma^{AB} p). \quad (\text{N.3})$$

The essential point is that a neural reduced model needs all three sectors. Pure gradient descent misses oscillation and phase transport. Pure Hamiltonian rotation misses irreversibility, gain control, and noise-driven switching. Pure diffusion misses learned structure and directed inference. A brain-scale runtime is therefore an excellent domain in which the gradient, curl, and noise sectors must coexist.

Table 70: One possible mapping from neural variables to the universal runtime grammar. The left column names the runtime object and the right a neural or cognitive interpretation of it. *One possible* is the operative phrase: the mapping is offered as an orientation for readers from neuroscience, the paper does not derive a neural quantum of information, and no result here depends on this table being the right assignment.

Runtime object	Neural or cognitive interpretation
$\mathcal{M}_{\text{neural}}$	Reduced manifold of activation, phase, memory, and action-readiness states.
$G_{AB}^{(\text{neural})}$	Effective information or control geometry; local distinguishability and transition cost.
$\mathcal{A}_A^{(\text{neural})}$	Recurrent phase transport, oscillatory coupling, and loop-induced nonconservative structure.
$\mathcal{F}_{\text{GFF}}^{(\text{neural})}$	Joint ledger of metabolic cost, mismatch cost, structural regularization, and reward or viability terms.
$\Sigma_{(\text{neural})}^{AB}$	Trial-to-trial variability, background fluctuations, and unresolved microstate contributions.
$\mathcal{O}_{\text{neural}}$	Readout dictionary from spikes or fields to behavior, report, or other recorded boundary observables.

N.2 DGI, GSSB, and GA in neural dynamics

The triad GRIP–GA acquires a particularly vivid form in neural systems. DGI says that the reduced generator should preserve its lawful content under admissible reparameterizations of coding variables, population coordinates, or gauge-like phase conventions. One should be free to rewrite the same dynamics in firing-rate variables, latent coordinates, or phase-amplitude coordinates without changing the physical predictions for observables in $\mathcal{O}_{\text{neural}}$. GSSB then states that concrete sensory context, task demands, developmental constraints, or injury can reduce the symmetry of the admissible generator, selecting some modes and suppressing others. GA is the observable outcome: a stable or metastable attractor family corresponding to a memory, percept, motor plan, or global action policy.

A practical way to write this is to separate the reduced drift into background-symmetric and context-selective parts:

$$\dot{z}^A = \mathcal{V}_{\text{sym}}^A(z) + \mathcal{V}_{\text{ctx}}^A(z; u) + \eta^A, \quad (\text{N.4})$$

where u denotes input or task context. DGI constrains how $\mathcal{V}_{\text{sym}}^A$ and $\mathcal{V}_{\text{ctx}}^A$ transform under admissible coordinate changes. GSSB is the statement that nonzero u , together with boundaries and long-run learning, reduces the effective symmetry of $\mathcal{V}^A$. GA is the resulting landscape of selected basins. The framework therefore predicts that neural function should not be understood as motion on a static graph only; it is motion on a graph whose effective symmetry class is continually being selected by boundary conditions.

N.3 Qi-HAL and the minimal adaptive closure

The main text introduced MePMoS as the minimal adaptive loop of Memory, Prediction, Motion, and Sensing. Here the usefulness of that decomposition can be made explicit. Let

$$m_{t+\Delta t} = \mathcal{M}(m_t, x_t, o_t), \quad (\text{N.5})$$

$$\hat{o}_{t+\Delta t} = \mathcal{P}(m_t, x_t, a_t), \quad (\text{N.6})$$

$$a_{t+\Delta t} = \mathcal{A}(m_t, x_t, \hat{o}_t, o_t), \quad (\text{N.7})$$

$$o_{t+\Delta t} = \mathcal{S}(a_t, \text{world}_t) + \nu_t. \quad (\text{N.8})$$

Equations (N.5)–(N.8) do not presuppose one specific neuroscience model. They state only that an adaptive open system capable of internal modeling must contain modules that play these four roles. Qi-HAL then adds higher-order structure on top of that loop: multiscale memory, a graph of attractor basins, phase-coded communication among scales, and switching operations that let the system pass from one effective mode to another without catastrophic loss of coherence. In other words, the MePMoS loop is the minimal closure, while the full holographic attractor lattice is the structured enlargement required for large-scale cognition.

N.4 Outer hologram, inner hologram, and evolving inner cosmos

The most delicate claim of the present program is not that the world contains measurement. That much was already established in GG-A5-GDOS. The more ambitious claim is that sufficiently complex open systems perform *internal* measurement as well. To state this carefully, define the outer-hologram map

$$o_t = \Pi_{\text{OH}}[\text{world}_t; a_t] + \nu_t, \quad (\text{N.9})$$

where o_t is the boundary record delivered by the world conditional on the system’s current action and sensory geometry. This is not passive reception; the action variable a_t explicitly enters the projection. Next define the inner-hologram dynamics

$$h_{t+\Delta t} = \mathcal{K}_{\text{IH}}(h_t, m_t, o_t, a_t), \quad (\text{N.10})$$

which updates an internal reconstructive state h_t on the basis of past internal state, memory, present sensory evidence, and current action. Finally define the slowly accumulated inner world-model

$$\mathcal{W}_{t+\Delta t} = \mathcal{U}(\mathcal{W}_t, h_t, r_t), \quad (\text{N.11})$$

with r_t the currently globally effective or reportable internal mode. The triple $(\Pi_{\text{OH}}, \mathcal{K}_{\text{IH}}, \mathcal{U})$ is the operational meaning of $\text{OH} + \text{IH} = \text{EIC}$.

Within this framework, “consciousness” is not introduced as a separate ontological substance. It is introduced as a regime of internal measurement in which some internal attractor state r_t becomes sufficiently stable, globally effective, and action-relevant that it functions as a boundary-visible state *for the system itself*. A simple operational criterion is

$$\tau_{\text{dwell}}(r_t) > \tau_{\text{cycle}}, \quad \Gamma_{\text{global}}(r_t) > \Gamma_c, \quad \mathcal{R}_{\text{action}}(r_t) > \mathcal{R}_c, \quad (\text{N.12})$$

where τ_{cycle} is the characteristic internal MePMoS cycle time, Γ_{global} is a measure of cross-module causal reach, and $\mathcal{R}_{\text{action}}$ is the degree to which the state constrains future selection. Equation (N.12)

is not claimed to be a final empirical law. It is a placeholder for the kind of measurable criterion the present framework demands.

N.5 Qi-CON and the spectrum of neural timescales

The neural domain is also one of the clearest places to test Qi-CON. Empirically, brain activity often exhibits scale-rich spectra rather than one clean resonance. In the present framework this is not an accidental superposition of unrelated processes. It is the spectral signature of an attractor lattice whose modes are arranged across nested timescales. The simplest predicted form is again a two-knee law,

$$S_{\text{neural}}(f) \sim \begin{cases} S_0, & f \ll f_{\text{global}}, \\ f^{-\alpha}, & f_{\text{global}} \ll f \ll f_{\text{local}}, \\ f^{-\beta}, & f \gg f_{\text{local}}, \end{cases} \quad (\text{N.13})$$

where f_{local} is associated with local recurrent loops and f_{global} with long-range integration cycles. The detailed exponents α and β are not universal numbers; they are effective summaries of architecture. The critical claim is instead that knee locations should correlate with explicit anatomical or functional hierarchy rather than floating independently of it.

Table 71: Candidate runtime criteria for reportable internal states. Five criteria run down the first column, with an operational reading of each and the reason it matters for the coupling of outer and inner holograms. They are labeled candidates because none is derived: the table’s function is to make a hard question tractable by proposing criteria that can be measured and argued with, and its third column says what each would settle if it held.

Criterion	Operational reading	Reason it matters for OH/IH/EIC
Metastable dwell time	State persists longer than the characteristic update cycle	Without dwell time the system cannot treat the state as an internal measurement record.
Broad causal reach	State influences multiple modules or scales	A purely local transient is not globally effective in the sense needed for report or coordinated action.
Action coupling	State constrains downstream policy or motor selection	Internal measurement matters only if it can reorganize future interaction with the world.
Reconstructive revision	State updates under new evidence without total collapse	The IH must behave like a revisable model, not a frozen label.
Energy-budget support	State maintenance consumes and justifies a measurable metabolic budget	The state must live on the same GFF ledger as the rest of the system.

N.6 Scope, caution, and falsifiability

The present appendix does not claim that every neural correlate of consciousness is already derivable from first principles. It makes a narrower and testable claim. If reportable conscious-like

states exist, then they should appear as a special class of globally effective attractor states in a multiscale internal-hologram architecture. If future experiments show that reportable states have no special dwell time, no broadened causal reach, no privileged relationship to action selection, and no identifiable role in the internal model, then the OH/IH/EIC proposal would be in serious trouble. A landmark theory paper must say that clearly.

O Ecological, social, and technological collective systems

The universal-runtime claim of GDOS is strongest, and most vulnerable, when it is exported beyond laboratory physics or biochemistry into collective systems. The danger here is obvious: one can too easily use scientific vocabulary metaphorically. The purpose of the present appendix is therefore to draw a strict boundary between legitimate export and category error. The theory does *not* claim that a society literally carries the same microphysical gauge charges as a single molecule. What it claims is that a collective system, once described as a finite open system with lawful ports, resources, memory, communication, and action channels, can be written on the same runtime grammar as any other reduced system. The reward for this discipline is that one may compare particles, cells, brains, and institutions without pretending they are identical. They are instead different realizations of one lawful open-system syntax (Dawkins, 1976, Friston, 2010, Kauffman, 1993, Shannon, 1948, Vicsek et al., 1995).

O.1 A collective System Instance Packet

Consider a collective system with population variables N_i , resource variables R_a , communication variables C_μ , memory or institutional variables M_λ , and action or policy variables A_κ . Define

$$z^A = (N_1, \dots, N_{N_N}, R_1, \dots, R_{N_R}, C_1, \dots, C_{N_C}, M_1, \dots, M_{N_M}, A_1, \dots, A_{N_A}). \quad (\text{O.1})$$

Then the runtime packet reads

$$\text{SIP}_{\text{coll}} = (\mathcal{M}_{\text{coll}}, z^A, G_{AB}^{(\text{coll})}, \mathcal{A}_A^{(\text{coll})}, \mathcal{F}_{\text{GFF}}^{(\text{coll})}, \Sigma_{(\text{coll})}^{AB}, \mathcal{O}_{\text{coll}}). \quad (\text{O.2})$$

The interpretation of each slot is now mesoscopic. $G_{AB}^{(\text{coll})}$ measures the local cost or sensitivity of displacing the collective state; $\mathcal{A}_A^{(\text{coll})}$ packages nonconservative institutional loops, delayed-control cycles, or technological feedbacks; $\mathcal{F}_{\text{GFF}}^{(\text{coll})}$ becomes the combined ledger of energetic or resource cost, entropy or uncertainty, structural maintenance, and coordination cost; $\Sigma_{(\text{coll})}^{AB}$ encodes unmodeled disturbances and heterogeneity; and $\mathcal{O}_{\text{coll}}$ is the explicit record dictionary: demographic rates, resource flows, communication delays, compliance, failure cascades, and resilience metrics.

O.2 Dynamic equations and the meaning of the four channels

At collective scale one may write the reduced equations schematically as

$$\partial_t N_i = B_i(z) - D_i(z) + \sum_j (F_{ji}(z) - F_{ij}(z)) + \eta_i, \quad (\text{O.3})$$

$$\partial_t R_a = \Phi_a^{\text{in}}(z) - \Phi_a^{\text{out}}(z) - C_a(z) + \zeta_a, \quad (\text{O.4})$$

$$\partial_t M_\lambda = \mathcal{T}_\lambda(M, C, A) - \mathcal{L}_\lambda(M) + \chi_\lambda, \quad (\text{O.5})$$

with analogous equations for communication and policy variables. The four-channel ledger then survives as a *bundled* description. J_E becomes the net energetic or resource-throughput channel required to maintain the system. J_e becomes the signaling or control-carrier channel, descended

from lower-level charged transport but now expressed as the effective communication or command throughput available at the chosen resolution. J_B becomes the structural-material maintenance channel: the one counting the throughput needed to preserve the physical or organizational body of the system. J_{Qi} becomes the effective information-throughput channel: update, coordination, discrimination, and model revision.

The export rule is therefore strict: the same symbols may be used across scales only when the observer has supplied a lawful bundling map from lower-level physical events to higher-level ports. Without such a bundling map, the notation would indeed be metaphorical and should be rejected. This insistence on explicit bundling is exactly how GDOS tries to remain scientific when exported to difficult domains.

Table 72: Permissible interpretation of the four current channels in collective systems. Each row is one channel, with a collective-scale reading. *Permissible* carries the fence: these readings are consistent with the framework but are not derived from it, and the paper does not claim that a social system locks its four currents through the same constant. The table is included because the question is asked often enough that leaving it unanswered invites a worse answer.

Channel	Permissible collective-scale interpretation
J_E	Resource or energy throughput needed to keep the system operating and to power adaptation.
J_e	Effective signaling or control throughput available to coordinate distant parts of the system.
J_B	Material and structural maintenance throughput preserving the body, infrastructure, or persistent organization.
J_{Qi}	Information-update throughput required for discrimination, learning, planning, and coordinated response.

O.3 Qi-EP for collective adaptation

The value of the runtime language becomes clearer when one considers adaptation. Collective systems often appear to optimize one quantity in one regime and another quantity in another regime. A firm may seem to minimize time to response during crisis, maximize throughput efficiency during routine operation, and minimize predictive mismatch during planning. An ecosystem may appear to maximize persistence of material cycling while minimizing vulnerability to shocks. From the perspective of Qi-EP, such plurality is expected. One may write a collective path functional

$$\mathcal{S}_{\text{coll}}[z] = \int dt \left[\frac{1}{2} g_{AB}(z) \dot{z}^A \dot{z}^B + \dot{z}^A \mathcal{A}_A^{(\text{coll})} + \Lambda_{\text{maint}}(z) + \Lambda_{\text{coord}}(z, \dot{z}) + \Lambda_{\text{forecast}}(z) \right], \quad (\text{O.6})$$

where the last three terms denote maintenance cost, coordination cost, and forecasting or model-mismatch cost. Short-response crises highlight the metric term and look PLT-like. Throughput-dominated steady operation highlights dissipation and looks MEPP-like. Long-horizon planning highlights mismatch reduction and looks FEP-like. The programmatic claim is again that one does not need separate constitutions for these cases; one needs only to understand which sector of the same path functional is dominant.

O.4 Qi-CON, heavy tails, and cascade structure

Collective systems are famous for burstiness, heavy tails, and cascade phenomena. The GDOS interpretation is not that such signatures are universally identical. It is that they are typical of multiscale attractor lattices with partial observation and delayed coupling. If communication, infrastructure, ecological turnover, and policy revision each operate on distinct timescales, then an observed aggregate signal should again display a knee-structured spectrum and possibly finite-window power laws. The generalized form is

$$S_{\text{coll}}(f) \sim \begin{cases} S_0, & f \ll f_{\text{inst}}, \\ f^{-\alpha}, & f_{\text{inst}} \ll f \ll f_{\text{local}}, \\ f^{-\beta}, & f \gg f_{\text{local}}, \end{cases} \quad (\text{O.7})$$

where f_{inst} is controlled by the slowest institutional or ecological revision cycle and f_{local} by the fastest local exchange or signaling loop. The same logic applies to burst-size distributions and waiting times. The important claim is not that every collective system obeys the same exponent. The important claim is that the existence and location of the scale window should be constrained by explicit architecture.

O.5 Where the framework can fail in collective domains

Because collective systems invite overinterpretation, the falsifiers must be stated sharply. The framework fails in this domain if the chosen bundling map is not explicit, if the proposed SIP packet does not improve prediction or diagnosis, if the current channels cannot be measured even approximately, or if the fitted scale-covariant laws bear no relationship to identifiable architecture. It also fails if the theory is used to smuggle teleology into a system that has no lawful memory, no meaningful action channel, and no persistent attractor structure. In such cases the honest conclusion would be that the export has exceeded its domain of validity.

O.6 Why the collective appendix still matters

Even with those cautions, the collective appendix is indispensable for a landmark paper. The reason is simple. If GDOS is truly a universal language for open systems, it must tell the reader what *sort* of export to attempt at the largest scales and what safeguards prevent abuse. A theory that stops at brains would still be a powerful theory. A theory that can say, in explicit runtime language, how ecological, social, and technological systems should be represented becomes something larger: a genuine bridge language across the five-story pagoda and beyond.

P Notation concordance and benchmark protocols

This last appendix records the editorial and methodological consequences of the present polishing round. It serves three purposes. First, it fixes notation that had previously appeared in parallel drafts with minor variants. Second, it states the line-breaking and layout conventions adopted in the source so that later revisions can be compiled with fewer overfull boxes and fewer accidental notational collisions. Third, it translates the entire manuscript into a practical benchmark program across the five-story pagoda. In that sense the appendix is both a concordance and a forward-looking protocol.

P.1 Notation concordance adopted in the polishing round

The notational rules below should be treated as authoritative for the current revision.

Table 73: Notation concordance. Each row names an object, gives the preferred notation for it and a comment. The table exists because the framework’s objects appear in several companion papers under slightly different symbols, and a reader moving between them needs one place that says which form this paper uses and why. It is a maintenance object rather than a result, and it is placed last for that reason.

Object	Preferred notation	Comment
Physics natural-unit current ratio	$J'_E : J'_e : J'_B :$ $J'_{Qi} = 1 : -1 :$ $-2/5 : -2/5 = 5 :$ $-5 : -2 : -2$	The order (E, e, B, Qi) is aligned with Bl-6.
Chemistry natural-unit current lock	$J_E/Q_B = J_e =$ $J_B = J_{Qi}$	This is the preferred form of the Arisaka Equation in chemistry-scale discussions.
Arisaka constant	$Q_B =$ $(2/5)(E_{\text{Ryd}}/K),$ $K = 35$	The symbol Q_B is preferred over earlier epsilon-based variants.
Information current	J_{Qi}	Use J_{Qi} , not shortened or ambiguous alternatives.
Triad order	$DGI \rightarrow GSSB \rightarrow$ GA	This order is conceptually important and should be preserved in prose and figures.
Qi-4 ordering	Qi-QUA, Qi-EP, Qi-CON, Qi-HAL	Use this canonical ordering unless a local pedagogical reason requires otherwise.
Dual holography identity	$OH + IH = EIC$	The equation is structural, not merely rhetorical.

P.2 Benchmark program

The five-story pagoda benchmark program (one candidate test per layer) is consolidated with the falsifiability program of Part XI: see Table 67 of Appendix L.1 for the integrated falsifiability matrix that supersedes earlier per-layer benchmark tables.

Part XIV

Common Objections and Responses

The strongest objections a hostile particle theorist, systems biologist, origin-of-life researcher, neuroscientist, and methodologist would each raise against the Qi-4 runtime — on the reality of the information sector, the Arisaka constant, the inevitability of life, the theory of consciousness, and the falsifiability of the whole — answered, with pointers into the paper

A construction that claims to extend one boundary law into a universal runtime for chemistry, life, brains, and society — and to price information as a gauge current with a definite energy quantum — invites the most demanding scrutiny; it would not be honest scholarship to leave the obvious objections unstated. We collect here the strongest adversarial readings of the Qi-4 thesis — the ones a particle theorist, a systems biologist, an origin-of-life researcher, a neuroscientist, and a skeptical methodologist would each raise on a careful reading — and answer each, with explicit pointers to the theorem, table, open problem, or companion paper that bears the weight. Where the honest answer is “not yet,” we say so and name the open item (in particular Open Problems O1–O9 and E1–E5 of §75). The objections are grouped by the discipline from which they are most naturally raised.

Q For the curious reader: the big questions, plainly answered

The twenty-five questions that follow are not a specialist’s. They are the ones a thoughtful reader without a physics degree actually carries into a paper like this one: whether information really costs anything and who pays, whether life is inevitable or merely possible, what a memory is when you look at it physically, whether consciousness has any business appearing in an equation, and what is left of biology if a cell turns out to be running the same machine as a star. Several of these have been settled in popular science in directions this paper does not take, and a reader is owed plain answers before the specialist objections begin. They are met first, and plainly.

The hardest of the twenty-five is the fourth — whether an account that reaches from particles to minds has explained anything about minds, or has merely used the word. The response concedes the difficulty rather than dissolving it, and says which part of it the paper can carry and which part it cannot. Where a claim rests on a step that is not closed, the response says so and names the step; where the honest answer is that the framework does not reach the question, it says that too.

Question 1 (Q-Qi4-1) (on the curious reader)

“Does information really have a price? That sounds like a metaphor.”

Response. It is not a metaphor, and the demonstration is sixty years old and has been measured. Landauer showed in 1961 that erasing one bit in a system at temperature T must dissipate at

least $kT \ln 2$ of energy. Not as an engineering limit but as a consequence of the second law: if erasure were cheaper one could build a machine extracting work from a single heat bath, which is forbidden. The bound has since been confirmed directly, in colloidal particles and single-electron devices. Bennett used the same insight to resolve Maxwell’s demon — the demon does not beat the second law because its memory must eventually be cleared, and clearing costs exactly what the sorting saved.

So information is already accountable in the same ledger as energy. What is new here is the exchange rate. Landauer’s price depends on the temperature of the bath, so it is a different number in a warm cell than in a cold detector. The claim in this paper is that there is a further, temperature-independent cost per event, fixed by the geometry rather than by the surroundings — roughly a sixth of an electron-volt per elementary informational event.

That is a much stronger claim than Landauer’s and correspondingly more exposed. It says every elementary event of this channel is billed at one rate, whether it happens in an enzyme, an electron transfer, or a memory write. One catalytic event whose four-way accounting misses by more than a fifth would end it.

Question 2 (Q-Qi4-2) (on the curious reader)

“You price a bit at about a sixth of an electron-volt, and the energy of the molecule life runs on is close to that. Isn’t that suspicious?”

Response. It should be, and the suspicion is the correct first response.

Here is the situation without decoration. The constant is assembled from the energy scale of the hydrogen atom and an integer fixed by a closure condition several papers earlier, for reasons that had nothing to do with biology and were fixed before biology was mentioned. It comes out at about 0.155 electron-volts per elementary event. The free energy released by the molecule that powers essentially every process in every living thing on this planet is within one and a half percent of that.

Either that is a coincidence or it is the most interesting sentence in the paper, and nothing about how the number was obtained settles which. What settles it is whether the same constant works elsewhere. The claim is not about one molecule; it is that every elementary event of this kind is billed at the same rate — an enzyme turnover, an electron transfer down a respiratory chain, a molecular motor step, a memory write in a neuron. Those are different systems with different chemistry, and the framework says one number prices them all.

So the honest position is: do not be persuaded by the agreement. Be persuaded, if at all, by whether the second, third and fourth cases come out at the same rate. Those measurements are ordinary biophysics rather than frontier physics, and they are the reason the claim is worth making at all.

Question 3 (Q-Qi4-3) (on the curious reader)

“Is life inevitable, or merely possible?”

Response. Neither in the form the question usually means, and the precise version is worth more than either.

What is claimed is a statement about a search rather than about a history. If a universe keeps

generating gradients, and if configurations capable of holding themselves together occupy more than a vanishing fraction of what is reachable, then such configurations are where the search ends up. That is inevitability in the ensemble sense — a statement about what happens given enough exploration, not a prediction about any particular planet.

What is emphatically not claimed: that any given world produces life. Nothing here says that of the Earth, and nothing here would be embarrassed by a sterile galaxy. Frequency is not the framework's subject and it does not compute one.

The condition doing the work is the one to watch, because it is where the claim could fail: that the viable region has positive measure and is repeatedly sampled. That is stated as a hypothesis rather than proved, it is flagged as such in the assessment, and quantifying it is one of the named open problems. A reader who wants the strongest version — life is written into physics — will not find it here, and should be suspicious of anyone offering it.

Question 4 (Q-Qi4-4) (on the curious reader)

“What is a memory, physically?”

Response. Not a place where something is stored, which is the picture worth giving up first. Storage suggests a location where a thing is put and from which it is later retrieved unchanged. Almost nothing about biological memory works that way, and the account here explains why it could not. What is held is not a set of items but a set of relations — differences between phases rather than the phases themselves. That is why the same content survives being carried on a completely different rhythm, so a chemical system, a nervous system and a machine can hold the same thing on hardware with nothing in common. It is why damage degrades a memory rather than deleting a list. And it is why “where is it stored” has no good answer, not because the answer is hidden but because the question presupposes the wrong kind of object.

What makes something a memory, on this account, is structural: a quantity that cannot be changed continuously. A count, once made, cannot be smoothly undone — there is no path from three windings to four — and that discreteness is what makes the record persist rather than drift away. Memory lives on the part of the structure that closes into a loop, and cost lives on the part that does not.

The consequence worth carrying away is that remembering is not a special biological trick. It is what any system does that has somewhere for a count to sit.

Question 5 (Q-Qi4-5) (on the curious reader)

“Is consciousness in your equations? And should I believe it if it is?”

Response. Something is, and it is much narrower than the word, and the narrowness is the point.

This is the hardest question in the set, so here is the boundary drawn as sharply as it can be. The framework defines a measurable quantity: how much of a system's internal model of its surroundings has been selected, stabilized, and made available to the system's own next cycle of action. That is a graded quantity, it applies to a bacterium and to a person by degree rather than by kind, and it can in principle be estimated for either. The paper gives it a name and studies it.

What it does not do — and this is where every account of this type either stays honest or does not — is claim that this quantity is what it is like to be the system. The question of why there is experience at all, rather than the same processing occurring in the dark, is untouched. The paper says so where the claim is made, and the caution is not a formality: a reader who takes the quantity for an explanation of experience has taken more than is offered, and the paper's own assessment marks the gap.

So: believe the quantity to whatever extent its predictions hold, which is a matter for measurement. Do not believe that anything here has explained consciousness. Nothing has, and the framework's willingness to say so is the reason to trust the rest of it.

Question 6 (Q-Qi4-6) (on the curious reader)

“The free-energy principle has a difficult reputation. Why should I accept it here?”

Response. Because here it is a consequence rather than a starting point, and that changes what can be said against it.

The reputation is deserved and worth stating. As usually presented, the principle says that any self-organizing system minimizes a quantity called variational free energy. Critics have argued for years that it is unfalsifiable — that the metric can be chosen after the fact, that the system boundary can be drawn to suit, and that the existence of an internal model is assumed rather than shown. Those are serious objections and they have not been fully answered by the principle's proponents.

What is different here is the direction of derivation. The principle is not assumed; it is obtained as one reading of a functional that was constructed from geometric data for other reasons. That has three consequences, and each addresses one of the standing objections. The metric stops being a modeling choice, because it comes from the structure. The system boundary stops being a hopeful assumption, because the projection defines it. And the existence of an internal model becomes a testable condition rather than a premise.

The caution belongs in the same breath. The inference reading fails exactly where the entropy readings fail — where circulation and gauge motion do work against the ledger — and the prospective-policy machinery that much of the literature builds on top is not claimed here at all.

Question 7 (Q-Qi4-7) (on the curious reader)

“If a cell runs the same machine as a galaxy, what is left of biology?”

Response. All of it. What is claimed is a shared grammar, not a shared subject, and confusing the two would make the claim both grander and emptier than it is.

The framework says that certain questions have the same form at every scale: what does this system spend, what does it hold on to, what would kill it, where does it stop. Those questions being shared does not make a cell a small galaxy any more than sharing a grammar makes a legal contract a love letter. Everything that makes biology biology — which molecules, which pathways, which contingencies of four billion years of history — is untouched and undetermined by anything here.

There is a real claim underneath and it is worth stating carefully. It is that biology and physics would stop being different *kinds* of explanation while remaining entirely different subjects. A

biologist would still need to know which enzyme, and no amount of geometry would supply it. What would change is that the answer to “why does this cost what it costs” would be the same kind of answer in both places.

The failure mode to watch for is the one this kind of claim usually falls into: explaining biology in a way that predicts nothing a biologist could check. The defense here is the constant — one number pricing an enzyme turnover, a motor step and a memory write, measurable in each case separately.

Question 8 (Q-Qi4-8) (on the curious reader)

“What is the physical difference between a rock and a bacterium?”

Response. A margin, and the fact that the answer is a number rather than a category is the substance of the claim.

The framework’s version is arithmetic. Something comes in; something is lost that cannot be recovered; something must be spent repairing what has gone wrong. If the first exceeds the other two on average, over a window long enough to include the repairs, the thing persists as itself. If it does not, for long enough, it stops. A rock has no repair term and does not need one; it persists by being hard rather than by maintaining itself. A bacterium maintains itself and would dissolve within minutes if it stopped.

There is no vital principle in that sentence and no threshold of complexity. There is a margin that can be positive or negative and that can in principle be measured for a protocell, a coral reef, a company or a state.

What quietly goes missing from the account is the moment at which something becomes alive, and the absence is the result rather than an omission. Anyone who has tried to say where the line falls — around a virus, a prion, a self-replicating crystal, an organelle that used to be somebody else — has found the line does not want to be drawn. The framework’s answer is that it is not there to be found. What is there is a plane with things scattered across it, and our words are names for regions of it rather than for kinds.

Question 9 (Q-Qi4-9) (on the curious reader)

“Is a thermostat conscious, then? Where is the line?”

Response. There is no line, and a thermostat sits so far down the scale that the question answers itself — but the reason there is no line is more interesting than either answer.

The framework’s quantity is graded. It measures how much of a system’s internal model has been selected, stabilized, and made available to the system’s own next cycle. A thermostat has a model of exactly one variable, holds it for as long as the bimetallic strip takes to bend, and has no next cycle in any meaningful sense. Its value is not zero as a matter of principle and it is negligible as a matter of measurement.

Refusing to draw a line is sometimes a way of avoiding a question, so it is worth saying why it is not here. Every attempt to draw one has failed in the same way: whatever criterion is proposed, cases are found just either side of it that differ by nothing anyone can articulate. That is the signature of a continuous quantity being forced into a binary. The framework takes the continuity seriously instead.

What this does not do is tell you whether a thermostat feels anything, and that limit should be repeated rather than glossed. A graded physical quantity says how much of a certain structure is present. Whether any amount of that structure is accompanied by experience is a question the framework does not reach and does not claim to.

Question 10 (Q-Qi4-10) (on the curious reader)

“Why do living things seem to fight entropy? Isn’t that against the second law?”

Response. They are not fighting it and never were, and the misunderstanding is one of the most durable in popular science.

The second law says the total entropy of an isolated system does not decrease. It says nothing about any part. A refrigerator makes its contents colder and more ordered every day by dumping more disorder into the kitchen than it removes from the cabinet. Living things do exactly this: they take in concentrated energy, use it, and radiate diffuse energy back out, and every ordered thing they build is paid for out of the difference. The Earth as a whole runs the same trade with the sun and deep space.

So the appearance of struggle is an artifact of drawing the boundary around the organism and forgetting the flow through it. There is no countercurrent and nothing being resisted.

What this framework contributes is a way of writing the trade with named terms rather than describing it qualitatively — intake, unrecoverable loss, repair cost — so that “staying alive” becomes an inequality that can be evaluated for a given system rather than a metaphor. Whether the particular inequality here is right is checkable. That the general picture is right has not been in doubt among physicists since Schrodinger’s lectures in 1943, and any account still describing life as defying the second law is eighty years out of date.

Question 11 (Q-Qi4-11) (on the curious reader)

“Does this explain the origin of life?”

Response. No, and the distinction between what it does and what that question asks is worth being exact about.

What the framework supplies is a set of conditions a system must satisfy to persist and to accumulate structure, together with a price for the informational events involved. That is a statement about what a candidate must look like, and it does constrain: it says what the accounting has to balance, and it prices the elementary events in units that a chemist could check.

What it does not supply is a route. Which molecules, in which order, on which mineral surface, under which conditions — none of that follows from anything here, and a framework that claimed it would be claiming to have done chemistry it has not done. The origin of life is a historical question about a particular planet, and constraints do not answer historical questions. There is one contribution worth naming because it is unusual. If the price of an informational event is fixed by geometry rather than by the local temperature, then the energy currency any emerging system must use is fixed too — not chosen by chemistry but imposed on it. That converts “why this molecule?” from an accident of early chemistry into a question with a

possible answer. Whether the answer is right is a laboratory matter, and the paper is careful to present it as a claim rather than a result.

Question 12 (Q-Qi4-12) (on the curious reader)

“Why should physics have anything to say about mind?”

Response. Because minds are made of things that obey physics, which is a weak premise almost nobody rejects, and because the specific structure at issue is not mental.

The move here is narrower than it sounds. The framework is about systems that maintain an internal model of what they cannot directly reach, and correct it against what arrives at their boundary. That description covers a bacterium holding a chemical trace of whether conditions have been improving, a rodent maintaining a map that survives its own turning, and a laboratory holding a theory of six dimensions and checking it against instruments that report four. The difference between these is depth, not kind.

That is a claim about information processing, and information processing in physical systems is a physical subject. It is not a claim about experience, and the paper keeps the two apart deliberately and repeatedly.

Where the objection does bite is on the risk of using the word to borrow credibility. A framework that reaches from quarks to minds has an obvious temptation to let the reach do the arguing. The protection is that the mind-adjacent claims are the most conditional in the paper and are marked as such, and that nothing else in the construction depends on them. A reader who deletes the last Part entirely still has the four pillars and every falsifier.

Question 13 (Q-Qi4-13) (on the curious reader)

“If systems maximize entropy production, why does anything look designed?”

Response. Because those are answers to different questions, and the apparent tension dissolves once the question each is answering is named.

The extremal principle picks out which of the available paths a system takes given its constraints. It does not choose the constraints. Design-like appearance comes from the constraints having been shaped, over long periods, by the survival of whatever satisfied them — and nothing about that process is in tension with a path being extremal at every instant. Water flowing downhill takes the steepest available route; that the riverbed has an intricate shape is a fact about erosion over millennia, not a violation of the flow.

There is a second point specific to this framework. The extremal readings hold only where circulation and gauge motion do no work against the ledger, and a system that has been shaped by selection is precisely one that has arranged for that not to be true — it maintains circulation deliberately. So the framework predicts that living systems will be exactly the systems where the simple extremal statements fail, which is a checkable claim and a satisfying one.

The appearance of design, on this account, is neither an illusion to be explained away nor evidence of a designer. It is what accumulated constraint looks like from outside, and the framework’s contribution is to say where the accumulation shows up in the equations.

Question 14 (Q-Qi4-14) (on the curious reader)**“Can a machine be alive by your criteria?”**

Response. By the persistence criterion, in principle yes, and the framework is deliberate about not flinching from that.

The criterion is arithmetic: intake exceeding unrecoverable loss plus repair cost, over a window long enough to include the repairs. Nothing in it mentions carbon, or reproduction, or having evolved. A machine that maintained its own components, repaired its own damage and secured its own energy would satisfy it, and no version of the criterion can be written to exclude such a thing without smuggling in a preference.

Two qualifications keep this from being a grand claim. First, no existing machine comes close; every one depends on a maintenance and supply chain that is not part of it, and drawing the boundary to include the chain changes what is being assessed. Second, the persistence criterion and the completeness index are different axes, and satisfying the first says nothing about the second — a thing can persist for a very long time while running almost none of the machinery that makes a system interesting.

There is a sharper point that the framework insists on and that is easy to miss. The criterion says nothing about whether a packet *deserves* to persist. The same arithmetic describes a healing wound and a tumor, a functioning institution and a captured one. What separates them is which boundary you chose to draw, and the framework is explicit that it does not choose for you.

Question 15 (Q-Qi4-15) (on the curious reader)**“Isn’t our universe just one bubble among countless others?”**

Response. It might be, and this paper is further from that question than others in the cohort, but it has one thing to add.

The multiverse is usually invoked to explain fine-tuning of the constants. There is a version of the fine-tuning list that this paper bears on directly: the entries about life. Those are often presented as the most striking — change this and there is no chemistry, change that and no stable molecules — and they are the entries where a selection effect is sharpest, because nothing incapable of persisting could be present to make a list.

What this framework adds is that the persistence condition is arithmetic rather than delicate. A system persists if intake exceeds loss plus repair. That is a much broader target than the fine-tuning discussion usually assumes, and if the viable region has positive measure rather than being a knife edge, the appearance of extraordinary luck is reduced without any ensemble being needed.

The condition is a hypothesis rather than a theorem, is flagged as one, and quantifying it is a named open problem. So this is a contribution to the argument rather than a resolution of it, and presenting it as more would be exactly the overreach the question is testing for.

Question 16 (Q-Qi4-16) (on the curious reader)**“The universe looks fine-tuned for life. Doesn’t that point to a designer?”**

Response. It points to something requiring explanation, and this paper contributes to the question from an unusual angle.

The cohort’s main answer is that the constants could not have been otherwise, which removes the inference’s premise. What this paper adds is about the other half of the fine-tuning list — the entries about life specifically.

Much of what strikes people as tuned for life is a statement about what it takes to persist, and persistence turns out to be an inequality with named terms rather than a narrow coincidence. If the viable region is not a knife edge, then a great deal of the apparent tuning is the ordinary observation that things which cannot hold themselves together are not around to be observed. That is a selection effect of a sharper kind than the usual anthropic argument, because it follows from the structure of persistence rather than from the requirements of chemistry.

This is not an argument against a designer and should not be recruited as one. Someone who believes on other grounds is untouched by anything here. What changes is that part of the fine-tuning list stops being evidence, because those entries could not have been observed to be otherwise by anything capable of making a list.

Question 17 (Q-Qi4-17) (on the curious reader)**“If everything is determined by geometry, is my future written? What happens to free will?”**

Response. Nothing in this paper touches your future, and this is the paper in the cohort with the most reason to be careful about the question.

The general answer holds: what is fixed is the rule book, not any history within it, and a rule book with no alternatives says nothing about which of an enormous number of permitted histories occurred.

The specific answer is that this framework describes systems that maintain an internal model and act on it, which is as close as physics gets to the machinery people have in mind when they ask about agency — and what it says is that such a system’s next action depends on a model it built, from evidence it gathered, through a loop nothing outside it computed. Whether that constitutes free will is a question about the phrase rather than about the physics, and philosophers who have thought about it longer than anyone here disagree.

What the framework does supply, and it is worth having, is a reason the question is not answerable by inspection. The state that determines what such a system does next is distributed across a model that only that system holds, and no shortcut computes it from outside. A universe can be lawful all the way down and still contain systems whose next move is not available in advance to anything but themselves.

Question 18 (Q-Qi4-18) (on the curious reader)**“Why is there something rather than nothing?”**

Response. The framework does not touch it, and saying so plainly is more use than a gesture. Everything here is conditional — given a structure of a certain kind, systems that persist must

satisfy this and pay that. Why there is a structure at all sits outside the construction entirely. That is not a gap further work closes; it is a different kind of question.

What this paper can contribute is narrower and is not nothing. If the price of an informational event is fixed, then the question of why anything persists has an answer of the same kind as the question of why anything falls: not a reason in the sense of a purpose, but a constraint that makes the alternative unavailable. That converts one member of the family of “why” questions into a member of the family of “what is forced” questions. It leaves the deepest member where it was.

The temptation, in a paper about emergence, is to say that something is more stable than nothing and therefore wins. That move is available, has defenders, and does not follow from anything here.

Question 19 (Q-Qi4-19) (on the curious reader)

“This explains an enormous amount from very little. Shouldn’t that make you suspicious?”

Response. It should and it does, and this paper carries the risk more heavily than the others in the cohort.

The general warning applies: Eddington derived the fine-structure constant from combinatorics and ended his career defending it. But the specific risk here is different and worse. A framework that reaches from gauge fields to metabolism to memory has an enormous surface on which a coincidence can be mistaken for a connection, and the reach itself is persuasive in a way that has nothing to do with evidence.

The protection is the constant. It is one number, fixed before biology entered the argument, and it prices every elementary informational event at the same rate. That is an unusually brittle commitment: an enzyme turnover, a motor step and a memory write are different systems with different chemistry, and if their four-way accounting does not come out at the same rate, the framework is finished regardless of how much it appears to explain.

So keep the suspicion, and aim it at the right place. Not at the reach, which is impressive and proves nothing, but at whether the second and third systems come out at the same price as the first. That is a bench measurement, and it is where this paper lives or dies.

Question 20 (Q-Qi4-20) (on the curious reader)

“Why should I take one person’s framework seriously against thousands of physicists and biologists?”

Response. Not because of who wrote it, and not despite it — the same error with opposite sign. The practical question is fair. The useful filter is exposure rather than headcount: a framework supported by thousands that forbids nothing measurable is a worse use of your attention than one supported by a single person that forbids specific things and names the measurement that decides. Community size tracks funding history and fashion and does not track correctness.

The cost of working outside a community is real. Errors are found faster where more people look, and a framework developed largely by one person carries a higher chance of a mistake that a hundred readers would have caught quickly. This paper’s subject makes the risk worse,

because it crosses fields, and the failure mode of cross-field work is a claim that looks solid to each field because each assumes the other checked it.

So do not take it on trust, and do not take the breadth as evidence. Take the constant, which needs no six-dimensional mathematics to evaluate: one number, claimed to price elementary informational events in systems as different as an enzyme and a neuron. Measure two of them. If the rates differ, the rest does not matter.

Question 21 (Q-Qi4-21) (on the curious reader)

“What would convince you that you are wrong?”

Response. Specific things, and this paper’s are unusually cheap to perform.

One elementary catalytic event whose four-way accounting misses the predicted rate by more than about a fifth would end the constant, and with it the informational pillar. A clean violation of the thermodynamic relations the framework says hold exactly in the lawful sector — the symmetry of a particular coupling matrix, a specific fluctuation relation — would break the variational pillar; these are bench measurements in cold atoms or mesoscopic devices, not frontier physics. A spectrum showing a single power law where the framework requires two turns, with the turn positions fixed by the system’s own slowest and fastest retained rates, would break the third. And a memory-bearing system that cannot be decomposed into the predicted normal form would break the fourth.

What would not convince: a failure of the consciousness-adjacent claims, which are already marked as the most conditional in the paper and on which nothing else depends; or a disagreement in a regime the framework says its readings do not cover — though that exception is itself constrained, since the framework states exactly where the readings fail rather than retreating to it after the fact.

Question 22 (Q-Qi4-22) (on the curious reader)

“Is the ‘self’ a physical quantity here?”

Response. A boundary is, and the difference matters more than it might seem.

What the framework can point at is a packet: a collection of processes with an accounting that closes, maintaining itself against loss, with a boundary defined by where the projection cuts rather than by anything anatomical. That is a physical object in the only sense the framework uses, and its persistence is computable.

What it cannot do is tell you which packet is you. The boundary can be drawn in many places — a cell, an organ, an organism, an organism plus its gut flora, a person plus their notebook — and the arithmetic works for each. The framework is explicit that it does not choose, and this is not evasion: the choice is what the word “self” is doing, and physics has no resource for making it.

There is a consequence worth sitting with. The same arithmetic describes a healing wound and a tumor, a functioning institution and a captured one. What distinguishes them is entirely which boundary you drew around what. If that seems unsatisfying, the dissatisfaction is pointing at something real — the framework has supplied the mathematics of persistence and nothing

whatever about whose persistence is worth wanting, and no amount of further physics will supply the second.

Question 23 (Q-Qi4-23) (on the curious reader)

“Does this say anything about whether there is life elsewhere?”

Response. Only in the weakest form, and it is worth being clear about how weak.

The framework says nothing about frequency. It does not compute how often life arises, because it does not compute histories — it computes constraints. A universe consistent with this framework could be teeming or could be empty, and neither would embarrass it. Anyone citing this program in support of an abundance estimate is overreaching.

There is one contribution that is not nothing. If the price of an informational event is fixed by geometry rather than by local chemistry, then the energy currency any living system must use is constrained wherever it arises. That is a claim about what alien biochemistry would have to look like at the level of the accounting, regardless of which molecules implement it — and it is checkable here before it is checkable anywhere else, because the same constraint applies to terrestrial life.

The honest summary is that this is a framework about what persistence requires, not about where it happens. The question you asked is a question about frequency and initial conditions, and the framework’s silence on it is structural rather than a gap awaiting further work.

Question 24 (Q-Qi4-24) (on the curious reader)

“Emergence is a word people use when they cannot explain something. Is that what is happening here?”

Response. It often is, and the criticism is fair often enough that a paper using the word owes an account of what it means by it.

The empty usage runs: the whole is more than the sum of its parts, therefore something new has appeared, therefore no further explanation is owed. That is not an explanation and it has been used to avoid a great many.

What is meant here is specific and checkable. A higher level is legitimate only if three things hold: its description closes approximately, so that it can be written without constant reference to the level below; it persists long enough to be worth describing; and it introduces genuinely new ports — ways of exchanging with its surroundings that the lower level does not have. Each of those is a condition that can fail, and where they fail the framework says there is no legitimate higher level.

The claim about levels is also weaker than it is usually made. Nothing here overrules anything below. What a higher level changes is never the law beneath it but the conditions under which that law is running — concentrations, boundaries, attention, institutions. The old dilemma, that a higher level must either reduce away entirely or acquire powers that break physics, turns out to be a false one, and the resolution is unglamorous: it was never a contest of causes but a difference of timescale.

Question 25 (Q-Qi4-25) (on the curious reader)

“I’m not a physicist and I’ve got this far. What is the one thing you want me to take away?”

Response. That information is not free, and that almost everything strange about living things follows from that one fact once you take it seriously.

For a century, energy has had a price in science and information has not. Information was the thing you could invoke when the physics ran out — the system “chose”, the molecule “recognized”, the network “learned” — and no bill ever came. That was not dishonesty. It was the absence of a rate. Without an exchange rate between energy and information, no statement about what a memory costs is checkable, and so no such statement is really a physical statement at all.

What this paper supplies is the rate. And the thing worth noticing is where it comes from, because it would have been much easier and much weaker to measure it. A number fitted to biochemistry would explain biochemistry and nothing else. This one is assembled from the energy scale of the hydrogen atom and an integer fixed several papers earlier for reasons that had nothing to do with life.

So take away this. If information really has a price fixed by geometry, then the familiar picture — in which physics describes the world and information theory describes our description of it — has the relationship backwards. There is one ledger, it has four columns, and the column we have been treating as bookkeeping is as physical as the other three. Nothing that remembers is getting anything for free.

R The particle theorist and mathematical physicist

Objection 1 (O-Qi4-1) (the gauge skeptic)

“Your whole edifice rests on $U(1)_{Q_i}$ being a genuine gauge sector, yet you concede it is ‘constitutionally optional’. If the constitution does not force it, none of your four pillars is a theorem of the framework — they are theorems of an unproven add-on.”

Response. We accept the logical point and state the status plainly: the four pillars are theorems of the $U(1)_{Q_i}$ -active class of admissible theories, not of the bare constitution. What the constitution forces is the four-component bulk $SU(5) \times U(1)_B \times U(1)_{Q_i} \times \text{gravity}$ as one consistent anomaly-free solution of BI-6; whether the $U(1)_{Q_i}$ factor is *empirically realized* is a separate, falsifiable question, recorded as Open Problem O9 and tested at high energy by the Kaxion search of Arisaka (2026, C1-Kaxion) (a predicted $8.1 \mu\text{eV}$ signature). So the pillars are conditional theorems with a sharp experimental gate, not unconditional ones — and we would rather say so than overclaim.

Objection 2 (O-Qi4-2) (the differential geometer)

“‘Information as a gauge current’ is a metaphor dressed as physics. Show me the conserved Noether current, the coupling constant, the boson, and one experimental signature — otherwise you have relabeled a bit-count as a charge.”

Response. The Noether current is the $U(1)_{Q_i}$ current J_{Q_i} of the bulk action imported from GG-A6-GRIP/GG-A2-GG6; the boson is the Kaxion, with the cohort-canonical mass $8.1 \mu\text{eV}$ ($\approx 1.95 \text{ GHz}$), and the signature is its direct detection in [Arisaka \(2026, C1-Kaxion\)](#). What we honestly concede is the *boundary* side: a turnkey protocol to count J_{Q_i} events in a cell or brain does not yet exist (Open Problem O2). So the high-energy current is a standard gauge current with a named boson; the open work is its low-energy metrology, which we flag rather than hide.

Objection 3 (O-Qi4-3) (the bioenergetics-minded reviewer)

“ $Q_B = (2/5)E_{Ry}/K$ with $K = 35$ is numerology: why the Rydberg, why divide by 35, why the 2/5? With that many discrete handles you can hit any biological energy.”

Response. The magnitude is no longer the soft spot. As now shown in Part III, the elementary quantum Q_B obeys $2 N_A Q_B = 30.0 \text{ kJ/mol}$ against the measured standard ATP-hydrolysis free energy $30.5 \pm 0.5 \text{ kJ/mol}$ — a 1.6% prediction, the integer multiplicity (one phosphoanhydride bond = two quanta) being forced, not fitted ([Arisaka \(2026, A11-Origin\)](#), GGA₇). The 2/5 is the Bl-6 ratio, E_{Ry} and $K = 35$ are imported from GG-A2-GG6/GG-A3-Bl6, and the chain has no continuous freedom. What *is* genuinely open (Open Problem O5) is the selection-rule *proof* that the per-step quantum is E_{Ry}/K rather than E_{Ry}/J — a discrete, not continuous, question — and we mark exactly that, no more.

Objection 4 (O-Qi4-4) (the skeptical reviewer)

“The Arisaka equation $J_e = J_B = J_{Q_i}$ is true by your choice of units: you insert a 5/2 in the definitions that exactly cancels the 2/5 in the ray. It is a tautology wearing the costume of a law.”

Response. The numerical equality after unit-normalization is indeed a convention, and we now say so at the point of display (Part III): the physical content of T-QUA.2 is not the equality but (i) the rank-one inflow structure — the four currents descend from a *single* one-dimensional source ray (Theorem 12.2), so they are not four independent quantities, and (ii) the existence of *one* conversion factor Q_B rather than four. The non-trivial, falsifiable claim is that one elementary boundary event deposits the fixed ratio $(1, -1, -2/5, -2/5)$ across all four sectors; that ratio is fixed by Bl-6, not chosen.

Objection 5 (O-Qi4-5) (the QED-precision reviewer)

“Your ‘two-knee $1/f$ ’ has no numbers: the slope is free and the knee positions are whatever the slowest and fastest rates happen to be. A bounded power law is the generic Lorentzian-superposition result known since the 1950s. What exactly is falsifiable?”

Response. We have sharpened the claim to what is actually derived (Part V, Theorem 42.1; F-Qi4-3): the structural prediction is the *co-occurrence of exactly two* scale-breaking knees bracketing a finite power-law band — not a single unbounded $1/f$ slope and not three or more

breaks — with the band slope tied to the conformal weight, and the two knees required to appear *jointly* with the other Qi-CON signatures (Lévy tails, fractal support). The knee positions are conceded to be the system’s slowest and fastest retained rates, not universal constants. That joint, bounded, two-break structure is the falsifier; a system showing a single unbroken slope, or breaks that do not co-occur with the other signatures, refutes it.

Objection 6 (O-Qi4-6) (the mathematical physicist)

“Theorem T-Qi4-9.3 (rank-one inflow) assumes a single BI-6 channel and concludes the source is rank one. That is ‘assume rank one, therefore rank one’. The only physics — that a second channel is forbidden — is offloaded to a citation.”

Response. Correct, and we do not claim otherwise: the uniqueness of the four-form X_4 (hence the single inflow channel) is *imported* from GG-A3-BI6, where it is the load-bearing result; within GG-A7-Qi4, T-Qi4-9.3 is the boundary consequence of that import, not an independent proof of channel-uniqueness. The honest reading is that the weight rests on GG-A3-BI6’s BI-6 uniqueness theorem, which we cite at the point of use; T-Qi4-9.3’s contribution is to carry that uniqueness to the four boundary currents.

Objection 7 (O-Qi4-7) (the string theorist)

“You import the GRIP cascade from GG-A6-GRIP as if it were proved, but its closure is conditional on GG-A6-GRIP’s open conjectures M10 and M11. Downstream you use it everywhere without re-flagging that.”

Response. Fixed, and you are right that it mattered. The conditional status is now re-flagged at every load-bearing use — the Qi-HAL universality theorem (Theorem 53.1), the pagoda factorization (Theorem 66.2), and the synthesis theorem (Theorem 77.1) each carry the inline caveat that the inherited GRIP-cascade closure is conditional on GG-A6-GRIP’s remaining residues — since the capacity theorem of Arisaka, 2026, P2-N3 resolved M10’s demand and its rigidity theorems closed the warped and relative refinements at class level, that means M11, now carried in spectral-flow form. Note the *other* import, the ground-attractor existence and uniqueness, does not depend on M10 or M11 at all — it is GG-A6-GRIP’s per-stratum corollary, carrying one stated assumption of its own (the moment-map identification) rather than either conjecture — so not everything downstream is conjectural, and the two conditionings should not be conflated.

Objection 8 (O-Qi4-8) (the operator-algebra specialist)

“Your SIP uniqueness proof is just Gelfand–Naimark plus Stinespring–Kraus: standard CPTP-channel identifiability. It says the packet data determine the dynamics — near-tautological for a well-posed semigroup — and cannot bear the weight of ‘one object, four faces’.”

Response. Agreed on the mathematics: the SIP completeness theorem (T-Qi4-8.2) uses standard results and proves exactly what you say, that two packets agreeing on all observables are gauge-equivalent. We do not rest the cross-domain universality on that lemma; the universality is a *constitutional* claim — the same GDOS boundary law and the same imported geometry apply at every scale — for which the lemma supplies only the well-posedness floor. Where the

prose over-read the lemma we have toned it (§9); the honest division is: standard machinery for well-posedness, constitution for universality.

Objection 9 (O-Qi4-9) (the falsifiability skeptic)

“Your headline falsifiability move is a ‘joint signature’ $\Xi_S = \sum w_a C_a$ with free weights and an undefined notion of ‘coherence’. A test with free knobs cannot be failed.”

Response. The weighted score Ξ_S is a heuristic summary, not the falsifier, and we should not have let it carry that rhetorical load. The actual falsifiers are the *named, weight-free* ones: the Arizaka bioenergetic quantum ($2Q_B$ vs measured ATP, currently +1.6%), the bounded two-knee co-occurrence (F-Qi4-3), the MePMoS/NHT/HAL/BFFT architecture in memory-bearing systems, and the high-energy $U(1)_{Qi}$ /Kaxion signature. Each can fail on its own measurement. The joint-coherence statement is the *additional* claim that these co-occur in a full-runtime system; it is suggestive, not the primary falsifier, and §75.3 now reads that way.

Objection 10 (O-Qi4-10) (the numerology-hunter)

“ $K = 35$ is asked to be a Diophantine identity, a logarithmic Planck-to-electroweak scale depth, and a linear count of atomic KK levels — three incompatible roles for one integer.”

Response. The Diophantine identity $K = LJ = M^2 + N^2 + L^2 = 35$ is GG-A3-BI6’s OGN-5 ledger; the scale depth and KK-level readings are the *same* K viewed through the renormalization-group dictionary of GG-A2-GG6, in which the fifth (scale) coordinate is the renormalization-group axis, so a logarithmic bulk depth and a level count are two descriptions of one ladder. We grant this is exactly the place a reader should demand the dictionary be explicit; the derivation lives in GG-A2-GG6/GG-A3-BI6, and GG-A7-Qi4 imports the single integer rather than re-deriving it. If GG-A2-GG6’s dictionary fails to map the depth onto the level count, that is a falsifiable defect of GG-A2-GG6, not a hidden tuning in GG-A7-Qi4.

Objection 11 (O-Qi4-11) (the consistent-with-versus-derived-from reviewer)

“Each pillar lists canonical results — Maxwell, Noether, Onsager, Friston — and asserts it ‘unifies’ them. Listing predecessors and claiming unification is not unification.”

Response. We have replaced ‘unifies these’ with ‘subsumes these as gauges or limiting cases of’ throughout the pillar prefaces, which is the accurate verb. For Qi-EP the reduction is genuine and explicit: PLA, MEPP, and FEP are shown to be three gauges of one projected functional $F_{GFF} = F_{GEP} + F_{CEP}$ (Theorem T-EP, the common-origin theorem), each recovered as a limit. For the others the honest claim is a common variational template, not a reduction with recovered coefficients, and the text now says so.

Objection 12 (O-Qi4-12) (the reductionist)

“Why drag a six-dimensional bulk into an information theory at all? It looks like cosmic overkill to motivate some bookkeeping about cells and brains.”

Response. The bulk is not introduced by GG-A7-Qi4; it is imported. The point of the program is that the *same* constitution that fixes the Standard-Model parameters (GG-A9-SM) and the

cosmic inventory (GG-A10-Cosmos) also fixes the open-system runtime — and that economy is the claim worth testing. If one wanted an information bookkeeping for cells alone, one would not need six dimensions; what six dimensions buy is that the bioenergetic quantum Q_B , the family number $N = 3$, and $\Omega_m = 1/\pi$ all descend from one integer ledger. That is either a deep unification or a coincidence, and the cohort is built to let the data decide.

S The systems biologist and biophysicist

Objection 13 (O-Qi4-13) (the bioenergeticist)

“ $Q_B = 0.155$ eV is not the energy of one ATP. Standard ATP hydrolysis is 0.32 eV and the in-vivo value is 0.5–0.66 eV. You are off by a factor of two to four and you call it your headline result.”

Response. You caught a real error, now corrected. Q_B is the *elementary* quantum; one ATP phosphoanhydride bond carries *two* of them, so the prediction is $2 N_A Q_B = 30.0$ kJ/mol against the measured *standard* free energy 30.5 ± 0.5 kJ/mol — a 1.6% match, not a factor-of-two miss (Part III; Arisaka (2026, A11-Origin)). The factor of two is the integer multiplicity, not a discrepancy. The in-vivo ~ 50 kJ/mol you cite is not a constant — it carries the mass-action term $RT \ln([ADP][P_i]/[ATP])$ — and even it is recovered as $(10/3)Q_B$ once the *E. coli* c_{10} proton stoichiometry is included.

Objection 14 (O-Qi4-14) (the biophysicist)

“What would you actually put on the axes to measure ‘eV per qubit’ in a cell? A qubit is not a biological observable.”

Response. The energetic side is measurable: Q_B is read off ATP-hydrolysis calorimetry as above. The *informational* side — counting J_{Qi} events independently of the energy, charge, and structural ledgers — is exactly what we flag as open (Open Problem O2): there is at present no turnkey instrument for it. So ‘eV per qubit’ is operationally half-grounded (the eV half) and half-promissory (the qubit half), and we label the promissory half rather than pretend it is settled.

Objection 15 (O-Qi4-15) (the stochastic-thermodynamicist)

“You call J_{Qi} conserved ‘like electric charge’, but your own balance law has generically nonzero information sources, and nobody can measure the flux. A conservation law you cannot measure and that is not conserved is an accounting convention.”

Response. The conservation is a boundary-current statement modulo named sources, not a closed-system identity, and the prose has been disciplined to say ‘conserved along its boundary and on appropriate time-scales’ rather than flatly ‘like charge’. The measurable proxy a systems biologist already uses — stochastic-thermodynamic mutual-information flow (Sagawa–Ueda-type bounds) — is the low-energy face of J_{Qi} , and connecting the two quantitatively is part of the O2 metrology program. Until that is done we do not claim J_{Qi} has the same operational standing as an ammeter current.

Objection 16 (O-Qi4-16) (the nonequilibrium-thermodynamicist)

“Your SIP and persistence theorems are autopoiesis and viability theory in new notation. What do they predict that Maturana–Varela, Aubin, or stochastic thermodynamics do not?”

Response. We concede the lineage openly — the persistence criterion (Theorem 58.1) is a Lyapunov/viability bound, and the SIP completeness theorem imports standard CPTP results. The added value is not a new persistence mechanism but a *constitutional embedding*: the same ledger that fixes Q_B and $N = 3$ also fixes the floor of the maintenance budget, so the persistence bound is tied to an integer-quantized energy scale rather than left as a free dissipation threshold. Whether that tie buys a novel, checkable number beyond the existing frameworks is the fair test, and we point to the bioenergetic quantum as the first such number.

Objection 17 (O-Qi4-17) (the molecular-machines reviewer)

“Your ‘universal’ decomposition into gradient, curl, noise, and gauge is just the Helmholtz–Hodge structure of *any* Fokker–Planck operator. It is true of every stochastic system, alive or not, so it forbids nothing.”

Response. Right that the decomposition is generic; the content is not in the decomposition but in the *ledger constraints* laid on it. Qi-QUA asserts that the budgets carried by the four sectors are quantized in units of Q_B and locked in the fixed BI-6 ratio — that is the part that can fail, and it is not a property of a generic Fokker–Planck operator. So the Hodge form is the universal canvas; the falsifiable claim is the Q_B -quantized, ratio-locked ledger painted on it.

Objection 18 (O-Qi4-18) (the ATP-synthase specialist)

“Your ‘worked example’ of ATP synthase is a citation plus the integer c_{10} . There is no closed bioenergetic audit.”

Response. It is now a closed audit at the level the framework claims (Part III): the standard bond energy is $2Q_B$; the per-proton chemiosmotic step $e\Delta p$ at the measured *E. coli* proton-motive force ≈ 0.155 V equals Q_B ; and the c_{10} ring with $H^+/\text{ATP} = 10/3$ recovers the in-vivo cost $(10/3)Q_B \approx 50$ kJ/mol. The deeper claim — that the rotor stoichiometry is an environment-tuned integer (c_8 mammals, c_{10} *E. coli*, c_{13} – c_{15} alkaliphiles/chloroplasts) selected by GRIP to match the available Δp — is forward-referenced to a dedicated forthcoming study, and labeled as such.

Objection 19 (O-Qi4-19) (the DNA biophysicist)

“You name-drop the 10.5 bp/turn helical repeat, writhe, and supercoiling, but no DNA observable is ever computed. Linking number is a real conserved quantity; you only place the words near each other.”

Response. Conceded: in GG-A7-Qi4 the DNA case is illustrative, not a closed derivation. The mapping we find suggestive — the linking-number identity $Lk = Tw + Wr$ as a structural-current conservation in the baryonic sector — is exactly that, suggestive, and the text now frames it as a motivating analogy rather than a worked result. We would rather under-claim DNA here than dress an analogy as a theorem.

Objection 20 (O-Qi4-20) (the chemotaxis modeler)

“Bacterial chemotaxis is quantitatively understood through Barkai–Leibler integral feedback. You recover none of that structure and then ‘predict’ the Levy walk, which is a sixty-year-old observation.”

Response. Fair on both counts. The framework supplies the runtime *grammar* (adaptive methylation as externalized HAL memory; run-and-tumble as a MePMoS loop), not the integral-feedback adaptation law, which it does not derive. The heavy-tailed search statistic is a retrodiction, and the verb has been softened from ‘predict’ to ‘is consistent with’. The honest status is organizational: we offer a common vocabulary across chemotaxis, neurons, and societies, not a replacement for the domain-specific kinetics.

Objection 21 (O-Qi4-21) (the falsification reviewer)

“The $\pm 20\%$ Qi-QUA kill-test is hollow: it tests a unit identity, its tolerance band excludes Q_B ’s own value, and you exempt every coarse-grained — i.e. every real — system as ‘overhead’.”

Response. The test is meant for *elementary* events — a single catalytic turnover — where the four ledgers can be measured independently and the equality is *not* guaranteed by unit choice (the unit choice fixes the ratio, not the per-event counts). Coarse-grained systems are excluded precisely because their overhead is built from sums of elementary events and belongs to a later bookkeeping layer; that is a scope statement, not an escape hatch. We grant the $\pm 20\%$ band needed recalibration to the elementary-event dispersion rather than the in-vivo ΔG spread, and Part III now states the elementary-event scope explicitly.

Objection 22 (O-Qi4-22) (the cell biologist)

“Concretely: what does this runtime let me compute that nonequilibrium thermodynamics does not?”

Response. Three numbers, in increasing order of sharpness: the quantized event budget $E = N_{\text{lock}} Q_B$ for an elementary catalytic cycle; the bounded two-knee spectral shape of a full-runtime system; and the bioenergetic anchor $2Q_B = 30.0$ kJ/mol, which nonequilibrium thermodynamics measures but does not predict. We are candid that much of the cross-domain content is organizational rather than predictive (Open Problems O1, O2), but the bioenergetic quantum is a genuine prediction a thermodynamicist would not write down.

Objection 23 (O-Qi4-23) (the origin-of-life skeptic)

“You invoke Schrödinger and say life needs ‘no new physics’, then introduce a new gauge sector and a new boson. Which is it?”

Response. Both, stated precisely: the *only* new physics is a single new $U(1)_{Q_i}$ gauge sector and its Kaxion boson; everything else — the open-system character of life, the information ledger, the persistence criterion — is reinterpretation of known physics. The Schrödinger framing is that life needs no new *biological* law beyond physics; it does need that one new gauge sector, which is a falsifiable high-energy commitment (Arisaka (2026, C1-Kaxion)). The text now owns this in one sentence rather than implying a free lunch.

Objection 24 (O-Qi4-24) (the proton-accounting reviewer)

“The thermodynamic H^+ /ATP for *E. coli* c_{10} is measured near 4.0, not your structural 10/3. Your ‘ $e \Delta p = Q_B$ ’ coincidence is then shaky.”

Response. Which is why the rigorous anchor is the *standard free energy*, $2Q_B = 30.0$ vs 30.5 kJ/mol, an assumption-free comparison that does not depend on the proton ratio at all. The per-proton picture ($e \Delta p \approx Q_B$ with the structural 10/3) is presented as suggestive mechanistic corroboration, and we explicitly note the thermodynamic-versus-structural H^+ /ATP distinction; reconciling them quantitatively, with the *c*-ring optimization across organisms, is the job of the forthcoming dedicated study, not of GG-A7-Qi4.

T The origin-of-life researcher**Objection 25 (O-Qi4-25) (the probabilist)**

“Your conditional-inevitability theorem is a Borel–Cantelli triviality: a recurrent process that hits a positive-measure set with bounded-below probability hits it almost surely. That is true of coin flips. The biology is smuggled into the hypotheses.”

Response. We agree it is a conditional reduction, and the title carries the word ‘conditional’ for that reason. The theorem’s job is not to prove inevitability *ex nihilo* but to isolate *exactly* what inevitability requires: a finite invariant measure, positive-measure viable windows, and bounded-below per-visit stabilization (Theorem 66.4). Naming those three conditions, and conceding (Open Problem O8) that we do not yet bound them, is the honest contribution; the abstract and summary have been retuned from ‘inevitable’ to ‘conditionally inevitable, given positive viable measure’.

Objection 26 (O-Qi4-26) (the prebiotic chemist)

“The whole result hangs on $\mu(W_{\text{life}}) > 0$, the rarity of a self-maintaining chemical recipe. That is the central unsolved problem, and you assert it as a sign condition.”

Response. Yes — and we say in Open Problem O8 that $\mu(W_{\text{life}})$ is ‘not provided in the present paper’. The theorem converts ‘is life inevitable?’ into ‘is the viable measure positive and repeatedly explored?’, which is progress only in that it localizes the hard question precisely. It does not answer it. We explicitly decline to claim the measure is large, or even nonzero, on present evidence.

Objection 27 (O-Qi4-27) (the major-transitions theorist)

“Why five floors? Maynard Smith and Szathmáry count about eight transitions. Your ‘particle to mind’ ladder is a curated reading of a generic adiabatic-elimination scheme.”

Response. Now stated outright in the pagoda section: the factorization theorem yields an *arbitrary-depth* tower wherever the floor-creation conditions hold, and the number five is an empirical reading, not a theorem output (Open Problem O7). A different coarse-graining — including a finer one matching the major-transitions count — satisfies the same existence

theorem. We keep the five-story figure as an interpretive aid and label it as interpretation, not derivation.

Objection 28 (O-Qi4-28) (the replicator-first/metabolism-first reviewer)

“‘Chemistry to life as GRIP-Out to GRIP-In’ relabels the origin of life without deriving autocatalysis, compartmentalization, or template replication. You define life as a positive maintenance margin and call the definition an origin.”

Response. Conceded: GG-A7-Qi4 states the *condition* a living system satisfies (the SeOS maintenance margin), not the *mechanism* by which chemistry crosses into it. We do not derive the error threshold, the concentration problem, or chirality, and we do not adjudicate metabolism-first versus replicator-first. The framework’s present contribution at this floor is descriptive vocabulary plus the bioenergetic anchor; the transition mechanism is named as future chemistry-floor work, not claimed.

Objection 29 (O-Qi4-29) (the stationarity reviewer)

“Borel–Cantelli needs a fixed invariant measure, but your selling point is that GRIP-Out events reorganize the landscape itself. You cannot have both a fixed measure and a measure that life keeps remaking.”

Response. A genuinely sharp point. The recurrence argument applies *within* a stratum — the fixed-landscape GRIP-In regime — where an invariant measure is well-defined; the cross-stratum GRIP-Out steps that remake the landscape are the separate, non-stationary question the theorem does not address. The honest consequence is that the infinite-trial idealization is the wrong limit for a finite, irreversibly evolving cosmos, and the physically relevant object is a finite-horizon hitting probability. We flag this as a limitation rather than papering over it; the finite-horizon estimate is future work.

Objection 30 (O-Qi4-30) (the compartmentalization reviewer)

“Your chemistry-floor equation has no compartment variable, yet your life-floor criterion presupposes a maintained membrane. You assume the very thing the field cannot yet explain.”

Response. True: compartmentalization enters GG-A7-Qi4 as an input boundary condition of the life floor, not as a consequence of the chemistry-floor dynamics. We do not derive self-organizing membrane closure from reaction–diffusion, and we should not be read as claiming to. It is named as an open boundary between floors — one of the places where the framework provides the destination’s description but not the journey’s mechanism.

Objection 31 (O-Qi4-31) (the astrobiologist)

“Your ‘conditional inevitability’ is compatible with life arising exactly once in the whole universe. The $N \rightarrow \infty$ limit is meaningless for a finite cosmos, so the slogan is empty.”

Response. Accepted, and this is precisely why the author’s own framing is ‘inevitable *under an Earth-like biosphere*’, not ‘inevitable everywhere’. With no bound on $\mu(W_{\text{life}})$ the theorem cannot distinguish ubiquitous from unique, and the abstract no longer implies it can. The

meaningful quantity is the expected number of origins per Hubble volume per Hubble time, which requires the measure estimates of Open Problem O8; we pose it and decline to answer it on present evidence.

Objection 32 (O-Qi4-32) (the no-novel-prediction reviewer)

“Every biological ‘test’ in your falsifiability matrix is either known phenomenology or an untestably general failure mode. Where is the risky origin-of-life prediction?”

Response. The risky, quantitative biological prediction is the bioenergetic one: $2Q_B = 30.0$ kJ/mol for the universal energy currency, already landing at +1.6%, with the sharper forthcoming claim that the rotor-ring integer tracks the biosphere’s proton-motive force. We concede the rest of the biological content is currently organizational or retrodictive (Open Problems O1, O2) and we do not dress retrodictions of $1/f$ noise or Levy walks as confirmations.

Objection 33 (O-Qi4-33) (the mind-window reviewer)

“Your second conclusion conditions mind-inevitability on a recurrent evolutionary exploration of the mind-window with bounded-below probability. That bakes in a directional drive to mind that evolution explicitly denies.”

Response. We agree there is no orthogenetic drive, and the conditioning $p_{\text{mind}} > 0$ is an assumption for which GG-A7-Qi4 supplies *no* selective mechanism. The theorem therefore says nothing about whether minds are favored or astronomically rare given life; it only says that *if* the mind-window is recurrently explored with positive stabilization probability, minds occur almost surely in the ensemble. That ‘if’ is doing visible work, and we leave it visible.

Objection 34 (O-Qi4-34) (the determinism reviewer)

“You headline ‘deterministically inevitable’ and then, mid-section, redefine determinism into harmlessness. That is a motte-and-bailey.”

Response. The body distinguishes the two meanings explicitly (§75.2): not that every micro-history is identical, but that higher-floor packets are attractor *classes* rather than freak accidents, in the ensemble sense. The strong trajectory word has been removed from the abstract and summary in favor of ‘ensemble-inevitable, conditional on positive viable measure’. The motte is now the only structure we defend.

Objection 35 (O-Qi4-35) (the inevitability-of-ATP reviewer)

“Is ATP-based life really inevitable, or is ATP just one chemical solution among many that you have retro-fitted to a constant?”

Response. The claim is conditional, not universal: *given* an Earth-like chemiosmotic biosphere with a proton-motive corridor, ATP is the ground geometric attractor of the bioenergetic stratum (GGA_7 , $2Q_B$), and the *c*-ring optimization shows the *same* attractor re-tuned across niches by an integer rotor count. A biosphere with no proton-motive chemistry could select a different currency, which is why we condition on the biosphere. The non-trivial content is that, conditioned on that chemistry, the energy quantum is fixed by geometry rather than by contingency — and that is what $2Q_B = 30.0$ kJ/mol tests.

U The neuroscientist

Objection 36 (O-Qi4-36) (the predictive-processing theorist)

“Your OH/IH/EIC loop is active inference plus the global workspace with new names. It states no experiment whose outcome differs from Friston, Tononi, or Llinás.”

Response. We concede EIC reproduces the predictive-processing and global-availability structure, and we do not claim a knockout of those theories. The framework’s distinctive content is the layer *beneath* them — a geometric constitution in which the free-energy update and the workspace access are boundary projections, with $U(1)_{Q_i}$ phase-coded memory and the bounded two-knee spectrum as the candidate differential levers. Where it cannot yet discriminate experimentally from active inference or IIT, we now say so rather than imply a settled theory of consciousness.

Objection 37 (O-Qi4-37) (the control theorist)

“Your ‘action-conditioned observability theorem’ is the Kalman observability rank condition from 1960, generalized to time-varying systems. Dignifying it with a cohort ID invites a desk reject.”

Response. Now credited at the theorem (§64.4): the rank condition is Kalman’s, and the active-sensing reading is Bajcsy’s and O’Regan–Noë’s. The contribution claimed is narrow and explicit — importing the time-varying observability criterion into the inner-hologram reconstruction, where the system’s own action sets the observation geometry. We do not present it as an original theorem, and the prose no longer does.

Objection 38 (O-Qi4-38) (the developmental neuroscientist)

“Your falsifier F-Qi4-5 — model degrades if action is restricted while passive input is preserved — is the Held–Hein kitten-carousel result from 1963. It was confirmed sixty years ago, so it is retrodiction, not a live frontier.”

Response. Conceded on the priority. F-Qi4-5 cites Held–Hein (and von Holst–Mittelstaedt’s reafference principle) and concedes the qualitative active-versus-passive asymmetry is established. The genuinely open, sharper prediction is the *quantitative* one — model error as a graded function of the observability-Gramian deficit, not a binary active/passive contrast — and that is what we now flag as the live test.

Objection 39 (O-Qi4-39) (the consciousness reviewer)

“Defining consciousness as ‘internal measurement of selected IH content’ buys an access criterion and then relabels the explanatory gap a ‘structural placeholder’. That renames the hard problem; it does not dissolve it.”

Response. Agreed, and now stated in those exact terms: the criterion the framework supplies is one of *access* consciousness — reportability and global availability of selected content, in the sense of Block (1995) — and the *phenomenal* hard problem is left explicitly open (Open Problem O3; the Conclusions now carry the access/phenomenal distinction). We claim a runtime condition for access-conscious content, not a solution to subjective experience, and we do not blur the two.

Objection 40 (O-Qi4-40) (the oscillations reviewer)

“Your $U(1)_{Qi}$ ‘memory carrier’ reduces to relative phase, and the holonomy is gauge-invariant precisely so that absolute phase is unobservable. A carrier built to be unobservable yields no neural test beyond ordinary phase coding.”

Response. The observable is indeed relative phase — phase precession, theta–gamma coupling — and we do not claim a measurement that separates a $U(1)_{Qi}$ -gauged memory from ordinary relative-phase coding (Open Problem O2). What the gauge structure adds is the constitutional *account* of why relative phase, corrected by holonomy, is the lawful carrier of a quantity whose absolute value is meaningless. That is an explanatory claim, not an extra neural observable, and the frequency-independent readout theorem (T-Qi4-12.9) is honestly just the gauge-invariance of relative-phase correlators.

Objection 41 (O-Qi4-41) (the aperiodic-spectrum reviewer)

“Your two-knee $1/f$ is indistinguishable from the standard FOOOF aperiodic-with-knee fit: free slope, fitted knees, generic f^{-2} roll-off. No knee in Hz is ever produced.”

Response. The knee positions are conceded to be fitted, as in FOOOF; the distinguishing content is structural and joint (F-Qi4-3): the prediction of *exactly two* co-occurring breaks bracketing a bounded band, the slope tied to the conformal weight, and the requirement that the spectral signature *co-occurs* with the Lévy and fractal diagnostics of the same selected branch. A single-knee aperiodic fit, or a two-knee fit whose breaks do not co-occur with the other Qi-CON signatures, is the discriminating failure. We do not claim a universal knee frequency in Hz, and the falsifiability text no longer does.

Objection 42 (O-Qi4-42) (the hippocampus reviewer)

“Your BFFT derives a dyadic 1:2 ladder, but the canonical hippocampal theta–gamma code nests at roughly 1:7. The data falsify your $\lambda = 2$.”

Response. Now conceded in the BFFT section: $\lambda = 2$ is the minimal idealization, and measured cross-frequency coupling is non-dyadic — theta–gamma at 1:5–1:8 (Lisman & Jensen (2013)), an integer $m \approx 7$. Proposition 51.1 already admits any integer $m > 1$; the dyadic case is the theoretical floor of the nesting hierarchy, not a claim that biology is binary. The empirical ladder selects the architecture-specific m , and we no longer present $\lambda = 2$ as a prediction.

Objection 43 (O-Qi4-43) (the energetics reviewer)

“ Q_B ‘per qubit’ applied to a neural memory write is unmotivated — a single spike costs 10^4 – 10^5 ATP for ion-gradient restoration, not one Q_B .”

Response. Correct, and the neural application is now scoped accordingly: Q_B is derived and validated at the *chemistry* (ATP-hydrolysis) scale, and GG-A7-Qi4 does *not* derive a neural per-bit quantum. A spike’s energy budget is dominated by Na^+/K^+ -ATPase, as you say; the statement that memory writes consume budget is qualitative (they do, via ATPase) and is not a claim of one Q_B per stored bit. We have removed any reading that prices neural information at Q_B directly.

Objection 44 (O-Qi4-44) (the IIT reviewer)

“Your graded continuum places a chemotactic bacterium at ‘proto-EIC’. That is the panpsychist move IIT is criticized for, smuggled past the reader in a table.”

Response. The continuum is a ladder of representational and architectural complexity, not a claim that every memory-bearing system is conscious — the non-claims table states exactly that. Part (d) of the universality theorem, which attaches ‘consciousness-grade’ to the top rung, is a definitional specialization, not a theorem corollary, and it is now read as such. We do not assert micro-experience for bacteria; we assert that the same runtime grammar runs at every rung, with consciousness reserved for the high-QCI, self-accessible end.

Objection 45 (O-Qi4-45) (the quantum-cognition skeptic)

“You write conscious readout as a POVM, $p(c) = \text{Tr}(\Pi_c \rho_{\text{IH}})$, with Luders post-selection — quantum-measurement formalism on a classical belief state. Either justify quantum neurons or drop the costume.”

Response. The instrument notation is the *general* measurement language, which reduces to ordinary Bayesian conditioning (marginalization over a partition of the belief simplex) precisely when the internal observable algebra is commutative — the relevant case for neural representations at brain temperature. We do not claim quantum neural states; the formalism is used because it specializes cleanly to the classical case while keeping the bulk-to-boundary projection language uniform. §64.7 now states the commutative reduction so the notation is not mistaken for a quantum-brain claim.

Objection 46 (O-Qi4-46) (the thalamocortical reviewer)

“You cite Llinás for ‘predictive coding’ — a mis-citation — and never engage his specific 40 Hz thalamocortical-resonance claim or IIT’s Φ -based dissociations.”

Response. The mis-attribution is corrected, and we concede the framework is presently *silent* on the signature predictions of the thalamocortical-resonance theory and of IIT — it neither reproduces nor contradicts the 40 Hz gating or the Φ dissociations. Engaging them on their own empirical commitments is named as future work rather than implied by a citation. ‘Relocating consciousness after the runtime’ is a statement about our exposition, not an empirical relation to those theories, and we no longer let it read otherwise.

Objection 47 (O-Qi4-47) (the localization reviewer)

“Your conscious-readout boundary ∂_{int} has no anatomical referent and no stated threshold — the appendix admits the criteria are placeholders. Any reportable state can be declared post hoc to have crossed it.”

Response. We accept that ∂_{int} is defined dynamically (global availability, dwell-time exceeding the cycle time, action-coupling) rather than anatomically, and that the thresholds are presently placeholders, as the appendix says. At this stage the consciousness construction is therefore a *research program* with an operational shape, not yet a fully falsifiable theory, and we mark it as such (Open Problem O3). Supplying order-of-magnitude thresholds — cycle times in milliseconds, a cross-module causal-reach metric — is the work that would promote it to a test.

V The methodologist and the working scientist

Objection 48 (O-Qi4-48) (the honest-disclosure auditor)

“Your abstract and conclusions speak as though the runtime is established, while your own triage lists thirteen of fourteen core problems as genuinely open. The front matter contradicts the body.”

Response. The front matter has been brought into line with the triage: ‘inevitable’ is now ‘conditionally inevitable’, the consciousness claim is scoped to access (Open Problem O3 governs), the two-knee falsifier states what is and is not predicted, and the imported GRIP closure is flagged conditional on the remaining conditioning (M11, spectral-flow form). The disciplined non-claims tables and the O1–O9/E1–E5 triage now govern the abstract, not the other way round. The candor was always in the body; it now reaches the summary.

Objection 49 (O-Qi4-49) (the single-falsifier reviewer)

“Strip away the joint signatures. What is the *one* experiment that, done tomorrow, could kill this?”

Response. The bioenergetic quantum. The framework forces $2Q_B = 30.0$ kJ/mol for the standard ATP-hydrolysis free energy; it sits at +1.6% of the measured value. A measured standard free energy disagreeing with $2Q_B$ beyond calorimetric precision, with no integer-multiplicity reinterpretation available, collapses the chemistry-floor anchor and with it the claim that the energy currency of life is geometrically fixed. That is a single number, already on the table, with a sharp tolerance.

Objection 50 (O-Qi4-50) (the cohort reviewer)

“Twenty-two interlocking papers look like a system engineered to absorb any refutation: knock one down and the others reabsorb it. Is the cohort falsifiable at all?”

Response. Each paper carries named, independent falsifiers — GG-A9-SM’s Standard-Model parameters, GG-A10-Cosmos’s $\Omega_m = 1/\pi$, Arisaka (2026, C1-Kaxion)’s $8.1 \mu\text{eV}$ axion, GG-A7-Qi4’s $2Q_B$ and two-knee signatures — and they share inputs, so a single measured failure propagates rather than being absorbed (a wrong $N = 3$, for instance, would impugn every paper that uses the OGN-5 ledger). The interlocking concentrates falsification risk; it does not diffuse it. The cohort map (§5.1) shows where each load-bearing input sits.

Objection 51 (O-Qi4-51) (the Occam reviewer)

“After importing GDOS, GRIP, the bulk, and the ledger, and conceding most floors to companion papers, what does GG-A7-Qi4 itself genuinely add?”

Response. Four pillar theorems (Qi-QUA, Qi-EP, Qi-CON, Qi-HAL), the internal-measurement EIC universality theorem, the Arisaka four-current lock, and the discovery that a memory-bearing system internalizes its own recoverable GRIP-Out (the Qi-HAL GRIP-loop) — the ten structural innovations of §74. We are explicit about what is imported (the constitution, the bulk, GDOS, the GRIP machinery) versus proved here, and the innovations are stated as the value added on top of those imports, not as a re-derivation of them.

Objection 52 (O-Qi4-52) (the standing-of-the-math reviewer)

“Much of your apparatus — GKLS semigroups, Hodge decomposition, Lyapunov bounds, Lorentzian-superposition spectra — is standard, competently relabeled. Where is the new mathematics?”

Response. We do not claim new mathematics; we claim a new *constitutional reading* of standard mathematics. The contribution is which standard objects are forced to appear, in which order, with which integer-quantized constants — not a new theorem in operator algebra or Morse theory. Where a result is textbook (CPTP identifiability, the two-knee Lorentzian superposition, the observability rank condition) we now cite the textbook source and claim only the application. The novelty is the architecture and the constants, not the lemmas.

Objection 53 (O-Qi4-53) (the working scientist)

“Why should I, with limited time, invest in reading and testing this rather than the framework I already use?”

Response. Because it makes a small number of sharp, risky, already-checkable commitments that your current framework does not: a geometrically fixed energy currency ($2Q_B = 30.0$ kJ/mol, at $+1.6\%$), a bounded two-knee spectral co-occurrence, a named memory architecture, and a high-energy $U(1)_{Q_i}$ axion at 8.1 μ eV — and it ties them to the Standard-Model and cosmological predictions of the same cohort through one integer ledger. If those commitments fail, the failures are diagnostic and localized; if they hold, disparate domains share one accounting. The paper is built to be wrong in specific, visible ways, which is the only honest reason to spend time on it.

Part XV

References

Cited works, each with a relevance note linking it to the present paper

Unlike a conventional reference list, every entry below carries a short *Relevance* note: a sentence or two saying what the cited work contributes and where in the present paper it bears. The aim is to turn the bibliography from a bare list of sources into a navigable companion — a reader can move through the references and see, at a glance, which theorem, table, or section each one underwrites. The notes are written for two readers at once. For the newcomer, they trace the intellectual lineage of the Qi-4 runtime: from the foundations of open-quantum-system dynamics, nonequilibrium thermodynamics, information geometry, and gauge theory, to the twenty-two-paper GG-Theory cohort in which the present paper supplies the four informational pillars (Qi-QUA, Qi-EP, Qi-CON, Qi-HAL) that run on top of the GDOS boundary law. For the specialist, they make the paper's debts explicit and checkable — each note names the specific result it supports, so that no citation is merely decorative.

The GG-Theory cohort papers (GG-A1-Constitution–GG-A11-Origin, together with the phenomenology P-series and the consequence C-series) are cited throughout under the canonical keys and titles fixed by the cross-paper coherence map of the program; the present paper's place in that cohort, and its dependence on the imported runtime constitution of GG-A6-GRIP, is set out in the entrance Part. The references are organized in two groups — first the author's GG-Theory papers (Section A below), then all other works (Section B), each in alphabetical order by first author. Every cited entry closes with a note recording the section in which it is first cited.

A. Papers by Katsushi Arisaka (the GG-Theory cohort)

Arisaka, K. (2026). *The GG-Constitution: Four Axioms and Seven Postulates of the GG-Theory Program, From Which GG-6, BI-6, OGN-5, and GDOS Descend as Theorems*. Arisaka-GG-A1-Constitution preprint, available at <https://preprints.arisaka-gg.org/a1-constitution/>. doi:10.5281/zenodo.22063867.

Relevance. The parent constitutional paper of the GG-Theory program. Establishes the four axioms AX-1 LC, AX-2 GGEP, AX-3 OGGEF, AX-4 MSMF, together with the seven postulates POS-1–POS-7 (GI, UEP, AF, CP, RG, OGN, MF) and the central Arisaka equation $J_E/Q_B = J_e = J_B = J_{Q_i}$ that ties the four boundary currents through one energy–information conversion factor. Every theorem in the present paper that names its constitutional source inline cites this paper at the point of invocation; the OGGEF axiom is the foundational dynamics from which the Qi-4 pillars and the EIC framework are derived as boundary consequences. First cited in §1.

Arisaka, K. (2026). *GG-Theory with GG-6: Geometry–Gauge Locking at the 6D Bulk (GG-6) at the Heart of Grand Geometric Theory (GG-Theory)*. Arisaka-GG-A2-GG6 preprint, available at <https://preprints.arisaka-gg.org/a2-gg6/>.

doi:10.5281/zenodo.22064298.

Relevance. The bulk-side foundational paper of the GG6 Trio. Establishes the 6D bulk metric $ds_6^2 = e^{-2r} \eta_{\mu\nu} dx^\mu dx^\nu + dr^2 + d\theta^2$, the canonical four-component gauge group $G_{\text{tot}} = \text{Spin}(1,5) \times \text{SU}(5) \times \text{U}(1)_B \times \text{U}(1)_{Qi}$, the magic-dimension ladder $\text{Spin}(1,5) \cong \text{SL}(2, \mathbb{H})$, and the OGN-5 Diophantine cascade $(M, N, L; J, K) = (1, 3, 5; 7, 35)$ with the dual-lock identity $K = LJ = M^2 + N^2 + L^2 = 35$. These are the ultraviolet inputs that the boundary GDOS dynamics of the present paper projects downward from; every Qi-EP variational statement that invokes the bulk X_4 four-form ultimately rests on the trace-ratio lock $|a/c| = |a/d| = 5/2$ established here. First cited in §1.

Arisaka, K. (2026). *Anomaly Freedom of BI-6: The Modified Bianchi Identity (BI-6) for the Six-Force, Six-Dimensional Bulk (GG-6)*. Arisaka-GG-A3-BI6 preprint, available at <https://preprints.arisaka-gg.org/a3-bi6/>.

doi:10.5281/zenodo.22064300.

Relevance. The anomaly-inflow companion of the GG6 Trio. Establishes BI-6 as the four-component Green–Schwarz-modified Bianchi identity $dH_3 = X_4$ on the 6D bulk, with the four-component coefficient vector $(a, b, c, d) = (1, -1, -2/5, -2/5)$ and the trace-ratio lock $|a/c| = |a/d| = 5/2$ pinned by the OGN-5 integer ledger. Section 8.7 (“BI-6 as the GGEP completion of GS: the geometrization reading”) is the canonical home of the supportive treatment of the 1984 Green–Schwarz construction that the present paper cites whenever the bulk anomaly-inflow structure is invoked. (This entry absorbs the legacy preprint code GD-2026-5, under which the BI-6 package circulated before the cohort renaming; older internal references to “Paper V” resolve here.) First cited in §1.

Arisaka, K. (2026). *From Euclid to GG-6: The Realization of Euclid’s Elements in Six Dimensions — GU-5, OGN-5, and a Parameter-Free Reconstruction of Modern Physics*. Arisaka-GG-A4-Euclid. Preprint, available at <https://preprints.arisaka-gg.org/a4-euclid/>.

doi:10.5281/zenodo.22064302.

Relevance. The historical companion to the GG6 Trio. Pedagogical commentary on the long chain from Euclid through Newton, Riemann, Maxwell, and Einstein to the GG-6 stage. We cite this paper whenever the historical positioning of the 6D bulk against its predecessors is invoked. First cited in §5.1.

Arisaka, K. (2026). *GDOS: Geometric Dynamics of Open Systems — 6D GG-6 Bulk Projected to 4D Boundary for Quantum Measurement, General Relativity, and QFT*. Arisaka-GG-GG-A5-GDOS preprint, available at <https://preprints.arisaka-gg.org/a5-gdos/>.

doi:10.5281/zenodo.22064304.

Relevance. The boundary-law foundational paper of the GDOS Trio, paired with the present paper (GG-A7-Qi4) and GG-A6-GRIP as the three formal pillars of the open-system half of the GG-Theory program. Derives the CPTP bulk-to-boundary projector $\mathcal{D}_{\text{OGGEP}} = \mathcal{R} \circ \text{Tr}_{\text{env}}$ from

AX-3 OGGEF; constructs the System Instance Packet (SIP) framework that the present paper uses to define the persistence class SeOS and the Qi-4 Completeness Index QCI; and establishes the boundary Lindblad-form dynamics that the Qi-4 pillars formalize on the information side. The Outer Hologram + Inner Hologram = Evolving Inner Cosmos (OH+IH=EIC) framework on the boundary, developed in detail in the present paper's Part IX, descends directly from the boundary-projection structure proved in this companion paper. First cited in §1.

Arisaka, K. (2026). *The GRIP on GDOS — GRIP: the Geometric Engine of Symmetry Breaking, from the Big Bang to Mind.* Arisaka-GG-A6-GRIP preprint, available at <https://preprints.arisaka-gg.org/a6-grip/>.
doi:10.5281/zenodo.22064306.

Relevance. Immediate predecessor to the present paper and the geometric half of the GDOS Trio. Formalizes the four Geometric Invariance Pillars (TGG, CCI, GFF, GRIP) as dual-layer principal theorems T1–T4 with named compatibility propositions $\mathfrak{C}_{\text{TGG}}$, $\mathfrak{C}_{\text{CCI}}$, $\mathfrak{C}_{\text{GFF}}$, $\mathfrak{C}_{\text{GRIP}}$. Derives the runtime constitution GRIP ($\text{DGI} \rightarrow \text{GSSB} \rightarrow \text{GA}$ cascade), the scalar ledger $F_{\text{GFF}} = F_{\text{GEP}} + F_{\text{CEP}}$, the SIP packet, and the Master Conservation Law $\nabla_\mu \mathcal{T}_{\text{TGG}}^{\mu\nu} = 0$. These are treated in the present paper as standing inputs for the parallel Qi-4 pillars (Qi-QUA, Qi-EP, Qi-CON, Qi-HAL), for the SeOS/QCI/EIC framework of Part IX, and for the five-story pagoda hierarchy of Part X. It is also the source of the runtime definitions adopted here: the Ground Geometric Attractor, established there as the unique global minimum on an admissible stratum and characterized by the four pillars in Theorem 62.1 of the present paper, and the two runtime regimes GRIP-In and GRIP-Out. First cited in §1.

Arisaka, K. (2026). *The Qi-4 Pillars — Qi-QUA, Qi-EP, Qi-CON, Qi-HAL: The Quantum-Information Backbone of GDOS.* Arisaka-GG-A7-Qi4 preprint (the present paper).
Preprint, available at <https://preprints.arisaka-gg.org/a7-qi4/>.
doi:10.5281/zenodo.22064308.

Relevance. The present paper. Formalizes the four information-theoretic Qi-4 pillars (Qi-QUA, Qi-EP, Qi-CON, Qi-HAL) descending from OGGEF on the memory-bearing $U(1)_{\text{Qi}}$ sector of the 6D bulk — the quantum-information half of the GDOS Trio. First cited in §5.1.

Arisaka, K. (2026). *Quantum Mechanics Derived from GG-Theory as Projection from 6D Bulk to 4D Boundary.* Arisaka-GG-A8-QM.
Preprint, available at <https://preprints.arisaka-gg.org/a8-qm/>.
doi:10.5281/zenodo.22064310.

Relevance. The first paper of the 4D-Physics Trio. Operationalizes T2 Operational Closure and T7 4D Quantum Mechanics from GDOS into a concept-by-concept reformulation of textbook quantum mechanics, with the central thesis that Einstein was right in 6D and Bohr was right in 4D, both within their respective domains. First cited in §5.1.

Arisaka, K. (2026). *The Standard Model as Theorem of GG-Theory: Its 28 Parameters as Observations of One Six-Dimensional Geometry.* Arisaka-GG-A9-SM preprint, available at <https://preprints.arisaka-gg.org/a9-sm/>.

[doi:10.5281/zenodo.22064312](https://doi.org/10.5281/zenodo.22064312).

Relevance. The 28-parameter codebook of the Standard Model derived from the OGN-5 integer ledger. Cited here as the downstream paper that uses the Qi-QUA rank-one inflow, the Arisaka constant Q_B , and the Qi-EP common-origin theorem to produce the closed-form Yukawa structure of the SM. First cited in §1.

Arisaka, K. (2026). *The Standard Cosmology as Theorem of GG-Theory: The Cosmic Inventory as Observations of One Six-Dimensional Geometry*. Arisaka-GG-A10-Cosmos preprint, available at <https://preprints.arisaka-gg.org/a10-cosmos/>.

[doi:10.5281/zenodo.22064314](https://doi.org/10.5281/zenodo.22064314).

Relevance. Derives the cosmological parameters $\Omega_m = 1/\pi$, $\Omega_\Lambda = 1 - 1/\pi$, $\Omega_b = 1/(2\pi^2)$, $\Omega_{DM} = 1/\pi - 1/(2\pi^2)$ from the GG-Constitution and the 6D bulk. Cited here as the companion paper carrying part of the resolution of the individual-environment realization of the pagoda hierarchy (open problem O4). First cited in §1.

Arisaka, K. (2026). *The Geometric Inevitability of Cosmos, Life, and Mind: The Self-Referential Closure of the GG-Theory Program*. Arisaka-GG-A11-Origin. Preprint, available at <https://preprints.arisaka-gg.org/a11-origin/>.

[doi:10.5281/zenodo.22064319](https://doi.org/10.5281/zenodo.22064319). No result of the present paper depends on it.

Relevance. The capstone of the program and the proof venue for the vocabulary adopted here: T-GGA-as-Ground-State, the cascade-wide generalization of the per-stratum ground-state result of GG-A6-GRIP whose four-pillar characterization is Theorem 62.1 of the present paper; the regime theorem T-GRIP-In-Lyapunov-Descend; and the Cascade-Self-Sustenance corollary whose microscopic mechanism is the Qi-HAL GRIP-loop (§52). Its eighteen-layer cascade is the full-length version of the five-story pagoda of Part X. First cited in §5.1.

Arisaka, K. (2026). *The GG-6 Lock Cluster: A Non-Spatial Kaluza–Klein Resonance near 3.85 TeV as the Standing LHC Falsifier of GG-Theory*. Arisaka-GG-P1-KK preprint, available at <https://preprints.arisaka-gg.org/p1-kk/>.

[doi:10.5281/zenodo.22064321](https://doi.org/10.5281/zenodo.22064321).

Relevance. The KK resonance-cluster prediction at $M_{\text{lock}} = (1/2)M_{\text{Pl}}e^{-35} \approx 3.85 \text{ TeV}$. Cited here as the downstream paper that uses the Qi-QUA Arisaka constant and the Qi-CON two-knee spectral theorem to predict the LHC-falsifiable mass-and-width structure. First cited in §5.1.

Arisaka, K. (2026). *Why Three Generations: A Three-Leg Derivation of $N_g = 3$ from BI-6, MSMF, and GDOS Morse Stability*. Arisaka-GG-P2-N3 preprint.

Preprint, available at <https://preprints.arisaka-gg.org/p2-n3/>.

[doi:10.5281/zenodo.22064323](https://doi.org/10.5281/zenodo.22064323).

Relevance. The second paper of the Predictions Branch: the family number derived as a squeeze between a CP-violation lower bound and a GDOS Morse/budget upper bound, with the conditioning declared by name. First cited in §5.1.

Arisaka, K. (2026). *The Electron as the Ground State of Nature: The Six Oldest Questions about the Electron, Closed and Fenced within GG-Theory*. Arisaka-GG-P3-Electron preprint.

Preprint, available at <https://preprints.arisaka-gg.org/p3-electron/>.
doi:10.5281/zenodo.22064325.

Relevance. The third paper of the Predictions Branch: identity, absolute stability, the two excited copies, the positron as conjugate readout, pointlikeness, and field response derived as six faces of ground-state-ness, with every fence declared; imports GG-P2-N3's $N_g = 3$ verbatim. First cited in §5.1.

Arisaka, K. (2026). *The Real Geometry Beneath the Complex Numbers of Physics: Positive Frequency is Holomorphy — from Cardano's Cubics to Twistor Space, via the Sixth Dimension of GG-Theory.* Arisaka-GG-P4-Real preprint.

Preprint, available at <https://preprints.arisaka-gg.org/p4-real/>.
doi:10.5281/zenodo.22064327.

Relevance. Foundations Branch: the origin-of- i member paper, deriving the single imaginary unit of boundary physics from the two real structures of the GG-A2-GG6 arena (the compact phase circle and the causal clock). For the present paper it is the origin of the boundary's complex structure whose phase the $U(1)_{Qi}$ memory carrier stores; cited at roster level. First cited in §5.1.

Arisaka, K. (2026). *The Mirrors of Matter: Why Every Particle Has a Twin, Why Nature's Mirrors Break, and Why the Universe Is Made of Matter.* Arisaka-GG-P5-CPT preprint.

Preprint, available at <https://preprints.arisaka-gg.org/p5-cpt/>.
doi:10.5281/zenodo.22064329.

Relevance. Foundations Branch: the CPT member paper, reading C, P, and T as orientation reversals of the three real bulk factors and carrying the matter–antimatter door; its mirrors and its law-level arrow of time are kin of the one-way information cascade that the Qi-4 runtime drives in a single direction; cited at roster level. First cited in §5.1.

Arisaka, K. (2026). *The Proton as the Ground Geometric Attractor of Matter: The Eleven Questions of the Proton, Closed and Fenced within GG-Theory.* Arisaka-GG-P6-Proton preprint.

Preprint, available at <https://preprints.arisaka-gg.org/p6-proton/>.
doi:10.5281/zenodo.22064331.

Relevance. Predictions Branch: the sister of GG-P3-Electron, running the ground-attractor method on the baryonic stratum. Its baryon number is a winding of the same compact circle whose phase this paper reads as memory, which makes it the sharpest available illustration of the difference between a quantity protected by a symmetry and one protected by a topological class; cited at roster level. First cited in §5.1.

Arisaka, K. (2026). *Symmetry, Conservation, and Geometry: Seven Steps from Noether, Bianchi, to Hopkins–Singer, and the Ledger the Open-System Turn Pays Out.* Arisaka-GG-W1-Symmetry preprint.

Preprint, available at <https://preprints.arisaka-gg.org/w1-symmetry/>.
doi:10.5281/zenodo.22064340.

Relevance. White-Paper Branch: the historical and conceptual survey of how conservation moved from symmetry to cohomology. It supplies the context in which Qi-QUA's four-current

locking is a statement about classes rather than about charges; cited at roster level. First cited in §5.1.

Arisaka, K. (2026). *The Circle That Remembers, the Interval That Forgets: GG-Theory and the Sixteen Problems of Particle Physics and Cosmology.* Arisaka-GG-W2-Holonomy preprint.

Preprint, available at <https://preprints.arisaka-gg.org/w2-holonomy/>.

doi:10.5281/zenodo.22064342.

Relevance. White-Paper Branch: the structural treatment of protection by topology. The memory that Qi-HAL formalizes is holonomy-class data on a compact fiber, so the ladder of protection mechanisms set out there is the classification this paper’s fourth pillar sits inside; cited at roster level. First cited in §5.1.

Arisaka, K. (2026). *The Kaxion: Cold Dark Matter as the Non-Spatial Compact-Angle Pseudoscalar of GG-6.* Arisaka-GG-C1-Kaxion preprint, available at <https://preprints.arisaka-gg.org/c1-kaxion/>.

doi:10.5281/zenodo.22064333.

Relevance. The dark-matter predictions companion of the GG-Theory program. Predicts a $\sim \mu\text{eV}$ Kaxion-like dark-matter candidate as the boundary realization of the $U(1)_{Q_i}$ sector at cosmological scales. Cited here as the high-energy companion paper that tracks the direct-detection program for the informational $U(1)_{Q_i}$ symmetry on which Qi-QUA, Qi-EP, Qi-CON, and Qi-HAL depend. First cited in §5.1.

Arisaka, K. (2026). *The Hubble Tension in the Light of GG-Theory: The Hubble Constant as One Reading of a Derived Expansion History $H(t)$.* Arisaka-GG-C2-Hubble preprint.

Preprint, available at <https://preprints.arisaka-gg.org/c2-hubble/>.

doi:10.5281/zenodo.22064336.

Relevance. The second paper of the Cosmology & Astrophysics Branch: reads the persistent discrepancy between early- and late-universe Hubble measurements as the two ends of one geometric corridor of the 6D bulk, rather than as a systematic error or new fluid. First cited in §5.1.

Arisaka, K. (2026). *Black-Hole Entropy under GG-Theory: The GDOS Reading of the Horizon as an OGGE Port — with the Exact Bekenstein–Hawking $A/4G_N$ Derived from Induced Gravity.* Arisaka-GG-C3-BH preprint.

Preprint, available at <https://preprints.arisaka-gg.org/c3-bh/>.

doi:10.5281/zenodo.22064338.

Relevance. The third paper of the Cosmology & Astrophysics Branch and the cohort’s strong-gravity anchor: the horizon treated as an OGGE port, with the exact $A/4$ coefficient derived. First cited in §5.1.

Arisaka, K. (2026). *One Functional, Three Principles: Maximum Entropy Production, Minimum Entropy Production, and the Free Energy Principle as Boundary Gauges of GDOS.* Arisaka-GG-S1-EP supplementary review preprint, available at <https://preprints.arisaka-gg.org/>.

doi:10.5281/zenodo.22190004.

Relevance. The extended companion of Part IV: the long-form histories of MinEP, MEPP, and FEP at full review depth, the complete proofs, the critical literatures assembled with their exact counterexamples, and the deployment of the three-gauge theorem with the grade of every step. Part IV is self-contained; the supplement is where every compressed statement made there can be checked at length. First cited in §20.

Arisaka, K. (2026). *Four Observations, One Exponent: Conformality and Scaling Laws, Lévy Flights, Fractality, and $1/f$ Spectra with Two Knees — the Measured Record, Audited*. Arisaka-GG-S2-Conformal supplementary review preprint, available at <https://preprints.arisaka-gg.org/.doi:10.5281/zenodo.22190070>.

Relevance. The measured record behind Part V: the four Qi-CON signatures audited against the literature at review depth, with the two-knee shape the claim that can fail. Cited where the scale certificate’s fixed-point language is read as continuation language. First cited in Part VIII.

B. Other references

Aguilera, M., Millidge, B., Tschantz, A., & Buckley, C. L. (2022). How particular is the physics of the free energy principle? *Physics of Life Reviews*, 40, 24–50.

Relevance. Exact analysis showing that the conditions for a statistical boundary to exist in the stationary density of a generic coupled system hold only on a narrow parameter subspace. The reason the inferential gauge here takes its system–environment split from the declared ports of the packet rather than hoping to find one. First cited in §24.

Amari, S.-I. (1998). Natural gradient works efficiently in learning. *Neural Computation*, 10, 251–276.

Relevance. Amari’s information-geometric formulation shows that optimization dynamics depend on the metric structure of probability/state manifolds. We cite it in the appendix to support POS-2’s claim that “extremal principles” on the boundary can be written as geometry-driven flows (e.g. natural gradients) when GDOS is expressed on statistical manifolds. First cited in §29.

Amari, S.-I., & Nagaoka, H. (2000). *Methods of Information Geometry*. American Mathematical Society.

Relevance. Information geometry provides the Fisher-metric viewpoint that underlies “natural gradient” flows. We cite it to motivate the GDOS interpretation of inference dynamics as geometric gradient descent on a scalar functional, and to connect GG-Theory’s GFF/GDOS language to established geometric statistics. First cited in §29.

Amari, S.-I., & Nagaoka, H. (2000). *Methods of Information Geometry*. American Mathematical Society.

Relevance. Information geometry provides a rigorous differential-geometric formulation of statistical models and inference via the Fisher metric and dual connections, serving as a mathematical precedent for writing GDOS as natural-gradient flow on a curved state manifold. First cited in §N.

Anderson, P. W. (1972). More is different. *Science*, 177, 393–396.

Relevance. Anderson’s argument that each level of complexity exhibits genuinely new organizing principles anticipates the pagoda hierarchy: the GDOS framework preserves level-specific emergence while denying level-specific lawlessness. First cited in §72.1.

Ashby, W. R. (1952). *Design for a Brain*. Chapman & Hall, London.

Relevance. Frames adaptive behavior as homeostasis of essential variables under ultrastable dynamics; the cybernetic root of the FEP’s identification of persistence with bounded surprisal. First cited in §24.

Atiyah, M. F. (1982). Convexity and commuting Hamiltonians. *Bulletin of the London Mathematical Society*, 14, 1–15.

Relevance. The perfection of moment-map Morse functions and the connectedness of their level sets, which extends Frankel’s uniqueness of the floor to the Morse–Bott case. One of the two theorems behind the uniqueness Theorem 62.1 imports. First cited in Part VIII.

Bajcsy, R. (1988). Active perception. *Proceedings of the IEEE*, 76, 966–1005.

Relevance. Founding statement of active perception: self-generated action improves the conditioning of the perceptual inverse problem. The lineage in which the action-conditioned observability of the IH sits. First cited in §64.4.

Ballhausen, B., Altendorf, K., & Deckers-Hebestreit, G. (2009). Constant c_{10} ring stoichiometry in the *Escherichia coli* ATP synthase analyzed by cross-linking. *Journal of Bacteriology*, 191, 2400–2404.

Relevance. Establishes the c_{10} ring stoichiometry of *E. coli* ATP synthase, providing a precise molecular integer that enters the event-budget accounting of the Arisaka Equation at the chemistry scale. First cited in §17.

Bauer, H. (1931). Dissipative dynamical systems I. *Proceedings of the National Academy of Sciences of the USA*, 17, 311–314.

Relevance. Proves that the equations of motion of a linearly damped system admit no variational formulation in the original variables: friction cannot simply be appended to a Lagrangian. This is the theorem that stopped the classical action tradition at the boundary of every open system, and it is why the trajectory gauge of this Part has to be derived rather than inherited. First cited in §21.

Belavin, A. A., Polyakov, A. M., & Zamolodchikov, A. B. (1984). *Infinite conformal symmetry in two-dimensional quantum field theory*. Nuclear Physics B, 241, 333–380.

Relevance. Founding paper of 2D conformal field theory; with Cardy 1984 establishes that scale invariance at criticality is automatically promoted to full conformal invariance. The structural ancestor of the conformality face of Qi-CON. First cited in Part V.

Bell, J. S. (1964). On the Einstein–Podolsky–Rosen paradox. *Physics*, 1, 195–200.

Relevance. Bell’s theorem clarifies that the measurement problem is not merely philosophical: operational predictions about correlations constrain the kinds of “hidden variable” completions

that remain possible. This motivates treating measurement as a constitutional issue rather than an interpretational afterthought. First cited in §1.2.

Benham, C. J. (1979). Torsional stress and local denaturation in supercoiled DNA. *Proceedings of the National Academy of Sciences of the USA*, 76, 3870–3874.

Relevance. Torsional stress and local denaturation in supercoiled DNA illustrate how geometric topology constrains molecular-scale dynamics, motivating the gauge-connection interpretation of structural currents in the chemistry-domain SIP. First cited in §17.

Bertini, L., De Sole, A., Gabrielli, D., Jona-Lasinio, G., & Landim, C. (2015). Macroscopic fluctuation theory. *Reviews of Modern Physics*, 87, 593–636.

Relevance. The comprehensive account of the quasi-potential and its Hamilton–Jacobi equation as the exact variational characterization of nonequilibrium steady states, of which the entropy-production readings are the near-equilibrium quadratic shadow. The independent statement of the order confinement derived here. First cited in §23.

Białynicki-Birula, A. (1973). Some theorems on actions of algebraic groups. *Annals of Mathematics*, 98, 480–497.

Relevance. The cell decomposition attached to the fixed points of a torus action; the stable set of each critical point of the descent on an admissible stratum is one of these cells, and “nothing can reach” is a statement about which cell is generic. First cited in Part VIII.

Biehl, M., Pollock, F. A., and Kanai, R. (2021). A technical critique of some parts of the free energy principle. *Entropy*, 23(3), 293.

Relevance. Proves the two blanket definitions inequivalent, the rewriting of expected flows invalid without unstated assumptions, and the Free Energy Lemma false at face value, with counterexamples. First cited in §24.

Block, N. (1995). On a confusion about a function of consciousness. *Behavioral and Brain Sciences*, 18, 227–247.

Relevance. The access/phenomenal distinction. The runtime criterion the present framework supplies is one of *access* consciousness; the *phenomenal* hard problem is left explicitly open. First cited in §75.2.

Bohr, N. (1913). On the constitution of atoms and molecules. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 26, 1–25.

Relevance. Bohr’s quantized atomic model establishes the discrete energy levels whose spacing ultimately descends from the Rydberg energy E_{Ryd} ; this energy scale reappears in the Arisaka Constant $Q_B = (2/5) E_{\text{Ryd}}/K$ as the bridge between physics-natural and chemistry-natural units. First cited in Part III.

Boltzmann, L. (1872). *Weitere Studien über das Wärmegleichgewicht unter Gasmolekülen*. Sitzungsberichte der Akademie der Wissenschaften zu Wien, 66, 275–370.

Relevance. The H -theorem and the first statistical-mechanical derivation of the second law from microscopic dynamics. The variational tradition that Qi-EP formalizes traces through this paper

to Onsager’s reciprocal relations and Prigogine’s minimum-entropy-production principle. First cited in Part [IV](#).

Boyer, P. D. (1979). *The binding-change mechanism for ATP synthesis*. in Membrane Bioenergetics (Lee et al., eds.), Addison-Wesley.

Relevance. The binding-change rotary mechanism of ATP synthase, providing the molecular machinery that realizes chemiosmotic ATP synthesis. First cited in Part [III](#).

Breuer, H.-P., & Petruccione, F. (2002). *The Theory of Open Quantum Systems*. Oxford University Press.

Relevance. This standard reference develops the mathematics of open quantum systems (reduced dynamics, master equations, Markovian limits). It anchors our claim that openness is not optional in realistic modeling, and it provides the mature language (CPTP maps and generators) that the GG-Constitution intends to elevate to first-class status. First cited in [§1.1](#).

Bruers, S., Maes, C., & Netočný, K. (2007). On the validity of entropy production principles for linear electrical circuits. *Journal of Statistical Physics*, 129, 725–740.

Relevance. The parity classification of the entropy-production principles by direct circuit analysis, and the identification of time-reversal-odd currents as the obstruction. Reached from below the resolution that the duality and failure theorems of Part [IV](#) derive from above. First cited in [§22](#).

Bruers, S. (2007). A discussion on maximum entropy production and information theory. *Journal of Physics A: Mathematical and Theoretical*, 40, 7441.

Relevance. Reconstructs both Dewar derivations, shows the 2003 machinery yields MinEP in an explicit model, and supplies the taxonomy of inequivalent MaxEP statements. First cited in [§23](#).

Bruineberg, J., Dolega, K., Dewhurst, J., and Baltieri, M. (2022). The Emperor’s New Markov Blankets. *Behavioral and Brain Sciences*, 45, e183.

Relevance. Distinguishes Pearl blankets (epistemic instruments within a model) from Friston blankets (reified physical boundaries), showing the mathematics does not underwrite the metaphysical reading. First cited in [§24](#).

Choi, M.-D. (1975). Completely positive linear maps on complex matrices. *Linear Algebra and Its Applications*, 10, 285–290.

Relevance. Choi’s theorem characterizes complete positivity, the mathematical core of physical admissibility for open-system dynamics and measurement instruments. AX-3 and POS-CP rely on complete positivity as the operational replacement for “no negative probabilities” across all open systems. First cited in [§9](#).

Colombo, M. and Wright, C. (2021). First principles in the life sciences: the free-energy principle, organicism, and mechanism. *Synthese*, 198(S14), 3463–3488.

Relevance. Reconstructs the FEP’s argument chain, exposes its stipulative cross-disciplinary identifications, and makes precise the point that a principle unfalsifiable by design cannot serve as an empirical law. First cited in [§24](#).

Conley, C. (1978). *Isolated Invariant Sets and the Morse Index*. CBMS Regional Conference Series in Mathematics 38, American Mathematical Society, Providence.

Relevance. The fundamental theorem of dynamical systems — every flow on a compact metric space has a complete Lyapunov function — and the Conley index, whose invariance under continuation is what the scale certificate’s fixed-point language means. First cited in §8.

Conant, R. C. and Ashby, W. R. (1970). Every good regulator of a system must be a model of that system. *International Journal of Systems Science*, 1(2), 89–97.

Relevance. The good-regulator theorem; the FEP’s claim that organisms are models of their environments is this result restated probabilistically. First cited in §24.

Crooks, G. E. (1999). Entropy production fluctuation theorem and the nonequilibrium work relation for free energy differences. *Physical Review E*, 60, 2721–2726.

Relevance. A cornerstone of stochastic thermodynamics connecting forward/backward path probabilities and nonequilibrium work. We cite it as a modern exemplar of how open-system probability flows carry rigorous, testable constraints—precisely the style enforced by POS-4 and POS-7. First cited in §17.

Da Costa, L., Friston, K., Heins, C., and Pavliotis, G. A. (2021). Bayesian mechanics for stationary processes. *Proceedings of the Royal Society A*, 477(2256), 20210518.

Relevance. The first mathematically careful Bayesian mechanics: for linear Gaussian processes with an imposed blanket, the synchronization map is exact and the expected internal flow is a genuine free-energy gradient flow. First cited in §24.

Dawkins, R. (1976). *The Selfish Gene*. Oxford University Press.

Relevance. The “selfish gene” perspective anticipates the information-centric view of biology; in the GDOS framework, information is a conserved gauge current (J_{Qi}), not merely a metaphor. First cited in §O.

Dayan, P., Hinton, G. E., Neal, R. M., and Zemel, R. S. (1995). The Helmholtz machine. *Neural Computation*, 7(5), 889–904.

Relevance. The entry of variational free energy into neural modeling: a recognition density trained to bound the log evidence of a generative model, mathematically identical to Friston’s F . First cited in §24.

de Groot, S. R. and Mazur, P. (1962). *Non-Equilibrium Thermodynamics*. North-Holland, Amsterdam.

Relevance. The standard treatise of linear irreversible thermodynamics; its Chapter V states the MinEP theorem with the constancy and boundary-condition hypotheses explicit. First cited in §22.

Dewar, R. C. (2003). Information theory explanation of the fluctuation theorem, maximum entropy production and self-organized criticality in non-equilibrium stationary states. *Journal of Physics A: Mathematical and General*, 36, 631–641.

- Relevance.** Dewar provides an information-theoretic route to maximum entropy production ideas in nonequilibrium systems. We cite it in POS-2 to represent the MEPP lineage and to clarify which claims are principled (variational) versus heuristic. First cited in §23.
- Dewar, R. C.** (2005). Maximum entropy production and the fluctuation theorem. *Journal of Physics A: Mathematical and General*, 38, L371–L381.
- Relevance.** Second derivation attempt, claiming Ziegler’s orthogonality far from equilibrium via a Gaussian step later shown to confine the result to linear response. First cited in §23.
- Dewar, R. C.** (2009). Maximum entropy production as an inference algorithm that translates physical assumptions into macroscopic predictions: don’t shoot the messenger. *Entropy*, 11, 931–944.
- Relevance.** The author’s concession that the Gaussian approximation is valid only close to equilibrium, and the retreat to MEP-as-inference. First cited in §23.
- Essex, C.** (1984). Radiation and the irreversible thermodynamics of climate. *Journal of the Atmospheric Sciences*, 41, 1985–1991.
- Relevance.** The structural objection that the dominant radiative entropy production is not extremized by the climate MEPP; the which-entropy-production problem. First cited in §23.
- Euler, L.** (1744). *Methodus inveniendi lineas curvas maximi minimive proprietate gaudentes*. Bousquet, Lausanne and Geneva (Additamentum II).
- Relevance.** The first rigorous statement of the least-action principle: for a single particle at fixed energy, the trajectory extremizes the abbreviated action. Euler’s appendix converted Maupertuis’s metaphysical economy into mathematics and founded the calculus of variations as the language of mechanics. First cited in §21.
- Fermat, P. de** (1662). Synthèse pour les réfractions (letter to M. Cureau de la Chambre). In *Œuvres de Fermat*, Vol. II (P. Tannery and C. Henry, Eds.), Gauthier-Villars, Paris, 1894, 457–463.
- Relevance.** The founding document of the least-time tradition: the sine law of refraction derived from the single postulate that light takes the path of least time. Every extremal principle reviewed in this Part descends, through the action tradition, from this calculation, and the objection it drew — why should nature optimize? — is the question the projected functional finally answers. First cited in §21.
- Feynman, R. P.** (1948). Space-time approach to non-relativistic quantum mechanics. *Reviews of Modern Physics*, 20, 367–387.
- Relevance.** The path integral makes the connection between action extremization and quantum amplitudes explicit (stationary phase in the classical limit). We cite it in POS-2 to link PLA to quantum mechanics within a single extremal framework. First cited in §21.
- Feynman, R. P., & Vernon, F. L.** (1963). The theory of a general quantum system interacting with a linear dissipative system. *Annals of Physics*, 24, 118–173.

Relevance. Introduces the influence functional for open quantum systems, a key tool for deriving reduced dynamics and decoherence. It is cited to connect OGGE/ GDOS-style open evolution to established path-integral machinery. First cited in §26.

Frankel, T. (1959). Fixed points and torsion on Kähler manifolds. *Annals of Mathematics*, 70, 1–8.

Relevance. The Morse indices of a moment-map component on a compact Kähler manifold are even, so the function is perfect and a connected stratum has one floor. The theorem behind the uniqueness Theorem 62.1 imports, and the reason the converse is a theorem on admissible strata. First cited in §8.

Freidlin, M. I., & Wentzell, A. D. (1998). *Random Perturbations of Dynamical Systems*, 2nd ed. Springer, New York.

Relevance. The quasipotential and the large-deviation rates between attractors of a small-noise flow; the quantities the halting certificate is written in on a stratum with several attractors, where an escape rate exists at all. First cited in Part VIII.

Friston, K., Kilner, J., and Harrison, L. (2006). A free energy principle for the brain. *Journal of Physiology—Paris*, 100(1–3), 70–87.

Relevance. First use of the phrase as a principle; introduces action as the second route to minimizing free energy, the germ of active inference. First cited in §24.

Friston, K., Daunizeau, J., Kilner, J., and Kiebel, S. J. (2010). Action and behavior: a free-energy formulation. *Biological Cybernetics*, 102(3), 227–260.

Relevance. The founding continuous-state active-inference paper: action minimizes the same free energy that perception minimizes, by changing the data rather than the beliefs. First cited in §24.

Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., and Pezzulo, G. (2017). Active inference: a process theory. *Neural Computation*, 29(1), 1–49.

Relevance. The discrete-state (POMDP) formulation with expected free energy over policies; the basis of essentially all empirical active-inference work. First cited in §24.

Friston, K., Da Costa, L., and Parr, T. (2021). Some interesting observations on the free energy principle. *Entropy*, 23(8), 1076.

Relevance. The authorized reply: accepts that the critique identifies implicit assumptions and heuristic proofs, and supplies a sparse-coupling ansatz restricting the principle’s antecedent conditions. First cited in §24.

Friston, K., Da Costa, L., Sajid, N., Heins, C., Ueltzhöffer, K., Pavliotis, G. A., and Parr, T. (2023). The free energy principle made simpler but not too simple. *Physics Reports*, 1024, 1–29.

Relevance. The path-integral reformulation that retreats from the nonequilibrium steady-state density factorization to a paths-of-least-action statement of the principle. First cited in §24.

Friston, K. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society B*, 360(1456), 815–836.

- Relevance.** First full statement of the brain-level program: hierarchical cortical processing as empirical-Bayes inference minimizing free energy. First cited in §24.
- Friston, K.** (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, 11, 127–138.
- Relevance.** Friston’s free-energy principle is a prominent modern attempt to unify perception, learning, and action by variational optimization. In the GDOS framework, FEP appears as one gauge of the universal Qi-EP extremal engine rather than as a standalone foundation. First cited in §1.1.
- Friston, K.** (2013). Life as we know it. *Journal of the Royal Society Interface*, 10(86), 20130475.
- Relevance.** The pivot to physics: random dynamical systems with Markov blankets at nonequilibrium steady state claimed to appear to perform Bayesian inference; origin of the contested derivations. First cited in §24.
- Friston, K.** (2019). A free energy principle for a particular physics. arXiv:1906.10184 (monograph, 148 pp.).
- Relevance.** The full physics-phase construction — particular partition, synchronization map, free-energy gradient flow — and the specific target of the technical critiques. First cited in §24.
- Fuchs, G. W., Cuppen, H. M., Ioppolo, S., Romanzin, C., Bisschop, S. E., Andersson, S., van Dishoeck, E. F., & Linnartz, H.** (2009). Hydrogenation reactions in interstellar CO ice analogues: A combined experimental/theoretical approach. *Astronomy & Astrophysics*, 505, 629–639.
- Relevance.** Demonstrates that complex molecular synthesis proceeds under driven nonequilibrium conditions even in interstellar environments, supporting the universality of the GDOS open-system framework beyond terrestrial chemistry. First cited in §17.
- Fuller, F. B.** (1971). The writhing number of a space curve. *Proceedings of the National Academy of Sciences of the USA*, 68, 815–819.
- Relevance.** The writhing number is a topological invariant of closed space curves; it enters the geometric characterization of DNA supercoiling and provides a concrete example of a curl-type conserved quantity in the chemistry-domain current ledger. First cited in §17.
- Glansdorff, P. and Prigogine, I.** (1954). Sur les propriétés différentielles de la production d’entropie. *Physica*, 20, 773–780.
- Relevance.** Introduces the decomposition $dP/dt = d_X P/dt + d_J P/dt$, the first step of the attempted extension beyond the linear regime. First cited in §22.
- Glansdorff, P. and Prigogine, I.** (1964). On a general evolution criterion in macroscopic physics. *Physica*, 30, 351–374.
- Relevance.** Proves $d_X P/dt \leq 0$ without linearity or reciprocity and establishes that $d_X P$ is not an exact differential — the precise sense in which MinEP admits no far-from-equilibrium generalization. First cited in §22.

Glansdorff, P. and Prigogine, I. (1971). *Thermodynamic Theory of Structure, Stability and Fluctuations*. Wiley-Interscience, London and New York.

Relevance. The mature statement of the program: MinEP confined to the linear regime, the local-potential methods, and the stability theory from which dissipative structures emerged. First cited in §22.

Goody, R. (2007). Maximum entropy production in climate theory. *Journal of the Atmospheric Sciences*, 64, 2735.

Relevance. Climate-side critique: the missing dynamical constraints, the steady-versus-chaotic distinction, and the planetary data assessment. First cited in §23.

Gorini, V., Kossakowski, A., & Sudarshan, E. C. G. (1976). Completely positive dynamical semigroups of N -level systems. *Journal of Mathematical Physics*, 17, 821–825.

Relevance. Along with Lindblad’s work, this establishes the generator form required for physical open-system quantum evolution. It supports POS-CP and the OGGE demand that open evolution be part of the foundational dictionary. First cited in §26.

Gray, C. G., & Taylor, E. F. (2007). When action is not least. *American Journal of Physics*, 75, 434–458.

Relevance. The definitive modern account of the fact that the realized path renders the action stationary rather than minimal, with the kinetic-focus criterion for when the extremum is a saddle. Cited to fix the correct universal statement — stationarity — which the projected functional of this Part and its three gauges inherit. First cited in §21.

Grinstein, G., and Linsker, R. (2007). Comments on a derivation and application of the “maximum entropy production” principle. *Journal of Physics A: Mathematical and Theoretical*, 40, 9717–9720.

Relevance. Identifies the error in the 2005 derivation and the unjustified self-organized-criticality claim of 2003; the pivotal critique. First cited in §23.

Haken, H. (1983). *Synergetics: An Introduction* (3rd ed.). Springer.

Relevance. Synergetics develops the slaving principle and order-parameter dynamics for self-organizing systems; these ideas are absorbed into the GSSB and GA stages of the master runtime pipeline. First cited in §72.1.

Hamilton, W. R. (1834). On a general method in dynamics. *Philosophical Transactions of the Royal Society of London*, 124, 247–308.

Relevance. Hamilton’s formulation is the modern backbone of the principle of least action and canonical dynamics. We cite it in POS-2 to emphasize that extremal principles historically reorganized mechanics and are a natural candidate for a unification engine. First cited in §21.

Hebb, D. O. (1949). *The Organization of Behavior: A Neuropsychological Theory*. Wiley, New York.

Relevance. Hebbian plasticity rule (“neurons that fire together wire together”), the founding model of synaptic learning dynamics. One of the four ancestral threads that Qi-HAL unifies. First cited in Part VI.

Held, R., & Hein, A. (1963). Movement-produced stimulation in the development of visually guided behavior. *Journal of Comparative and Physiological Psychology*, 56, 872–876.

Relevance. The kitten-carousel experiment: actively moving animals develop visually guided behavior while passively moved animals with identical optic flow do not — the classic confirmation of the active-versus-passive asymmetry the action-conditioned observability theorem predicts. First cited in §75.3.

von Helmholtz, H. (1856–1867). *Handbuch der physiologischen Optik*. Voss, Leipzig; three volumes, completed 1867.

Relevance. Introduces unconscious inference, the doctrine that percepts are conclusions drawn from sensory premises and prior experience; the conceptual ancestor claimed by predictive coding and the free energy principle alike. First cited in §24.

Helmholtz, H. von (1868). Zur Theorie der stationären Ströme in reibenden Flüssigkeiten. *Verhandlungen des naturhistorisch-medizinischen Vereins zu Heidelberg*, 5, 1–7.

Relevance. The minimum dissipation theorem of Stokes flow, establishing the pattern that the minimum property belongs to the linear regime alone. First cited in §22.

Higgs, P. W. (1964). Broken symmetries and the masses of gauge bosons. *Physical Review Letters*, 13, 508–509.

Relevance. Foundational paper on gauge-boson mass generation. GG-6 reproduces the mass relations from holonomy. First cited in §7.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences of the USA*, 79, 2554–2558.

Relevance. Hopfield networks provide the canonical model of attractor-based associative memory; in the Qi-HAL framework, the Holographic Attractor Lattice generalizes the Hopfield picture to a metastable landscape of selected attractors under open-system dynamics. First cited in Part VI.

Jarzynski, C. (1997). Nonequilibrium equality for free-energy differences. *Physical Review Letters*, 78, 2690–2693.

Relevance. Jarzynski’s equality shows that free-energy differences can be inferred from nonequilibrium work statistics. It supports the UEP/GDOS viewpoint that variational/ledger principles remain meaningful and falsifiable far from equilibrium. First cited in §17.

Jaynes, E. T. (1957). Information theory and statistical mechanics. *Physical Review*, 106, 620–630.

Relevance. Introduces the maximum-entropy (MaxEnt) principle as a constructive variational rule for statistical mechanics. We cite it in POS-2 to connect thermodynamic extremal principles to inference and to motivate treating free-energy/entropy functionals as legitimate variational objects when paired with operational constraints. First cited in Part IV.

Jaynes, E. T. (1980). The minimum entropy production principle. *Annual Review of Physical Chemistry*, 31, 579–601.

Relevance. The scope critique that distinguishes Onsager’s principle from Prigogine’s and exhibits a two-temperature network in which the minimum principle fails. In the present framework that network is a case with surviving curl-sector work, and the failure is derived rather than cataloged. First cited in §22.

Jiang, W., Hermolin, J., & Fillingame, R. H. (2001). The preferred stoichiometry of c subunits in the rotary motor sector of *Escherichia coli* ATP synthase is 10. *Proceedings of the National Academy of Sciences of the USA*, 98, 4966–4971.

Relevance. Confirms the preferred c_{10} stoichiometry of the ATP synthase rotor in *E. coli*, a key experimental integer used in the event-budget verification of the Arizaka Equation. First cited in §17.

Kalman, R. E. (1960). On the general theory of control systems. *Proc. First IFAC Congress*, Moscow, 481–492.

Relevance. The classical observability criterion: a linear (time-varying) system’s initial state is recoverable iff the observability matrix has full column rank. The action-conditioned observability theorem behind the inner hologram is this criterion applied to internal-model reconstruction. First cited in §64.4.

Kauffman, S. (1993). *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press.

Relevance. A foundational systems-biology perspective emphasizing attractors, networks, and self-organization in evolution, connecting the attractor and phase-selection language to established biological dynamical-systems frameworks.

First cited in §M.

Kirchhoff, G. (1848). Ueber die Anwendbarkeit der Formeln für die Intensitäten der galvanischen Ströme in einem Systeme linearer Leiter auf Systeme, die zum Theil aus nicht linearen Leitern bestehen. *Annalen der Physik und Chemie*, 75, 189–205.

Relevance. The earliest minimum-dissipation theorem for a driven steady state — currents in a linear resistor network minimize the Joule heat production — and the benchmark against which Jaynes’s counterexample is measured. First cited in §22.

Klein, M. J. and Meijer, P. H. E. (1954). Principle of Minimum Entropy Production. *Physical Review*, 96, 250–255.

Relevance. First statistical-mechanical verification of Prigogine’s theorem, whose first-order-only agreement is the earliest quantitative marker of MinEP’s status as a near-equilibrium approximation. First cited in §22.

Kossakowski, A., Frigerio, A., Gorini, V., & Verri, M. (1977). Quantum detailed balance and KMS condition. *Communications in Mathematical Physics*, 57, 97–110.

Relevance. The quantum detailed-balance condition, satisfied by a weak-coupling generator whose bath is in a KMS state; the hypothesis under which the relative-entropy descent of Spohn

(1978) is a free-energy descent, and the hypothesis of GG-A6-GRIP’s Theorem T3.2. First cited in §8.

Kraus, K. (1983). *States, Effects, and Operations*. Springer.

Relevance. Kraus’ monograph systematizes the operational formulation of quantum measurement and dynamics via completely positive maps. AX-3 adopts this language constitutionally by requiring that measurement be defined in terms of instruments and that open evolution be CPTP. First cited in §9.

Lagrange, J.-L. (1788). *Mécanique analytique*. Veuve Desaint, Paris.

Relevance. The construction of the whole of mechanics on the variational calculus, with neither figures nor Maupertuis’s theology. Lagrange’s treatise is the point at which the action principle stopped being a curiosity about light and became the standard form of physical law. First cited in §21.

Landauer, R. (1961). Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5, 183–191.

Relevance. Landauer’s principle connects information processing with thermodynamic cost. We cite it in the undergrad-intuition appendix to emphasize that “information” is physical in open systems, reinforcing GDOS’s treatment of boundary dynamics as energetics + entropy + information bookkeeping under constitutional constraints. First cited in §1.1.

Landauer, R. (1975). Inadequacy of entropy and entropy derivatives in characterizing the steady state. *Physical Review A*, 12, 636–638.

Relevance. The blowtorch theorem: steady states of inhomogeneously heated bistable systems depend on kinetic details invisible to the entire entropy family of functionals. First cited in §22.

Landauer, R. (1975). Stability and entropy production in electrical circuits. *Journal of Statistical Physics*, 13, 1–16.

Relevance. Exhibits the failure of MinEP arbitrarily close to equilibrium in circuits with inductors — later diagnosed as a violation of the tacit time-reversal-parity hypothesis. First cited in §22.

Lane, N. (2015). *The Vital Question: Why Is Life the Way It Is?* W. W. Norton & Company.

Relevance. Lane’s emphasis on energy flow in the origin of life aligns closely with the GDOS GFF ledger approach. First cited in §M.

Lévy, P. (1937). *Théorie de l’addition des variables aléatoires*. Gauthier-Villars, Paris.

Relevance. Founding text on stable distributions and the Lévy stable laws. Structural ancestor of the heavy-tailed statistics of Qi-CON. First cited in Part V.

Lindblad, G. (1975). Completely positive maps and entropy inequalities. *Communications in Mathematical Physics*, 40, 147–151.

Relevance. Relative entropy is non-increasing under any CPTP map, so the relative entropy to a stationary state descends under any CPTP semigroup that has one — the level at which Theorem 8.2 sits. First cited in §8.

Lindblad, G. (1976). On the generators of quantum dynamical semigroups. *Communications in Mathematical Physics*, 48, 119–130.

Relevance. Lindblad’s theorem provides the canonical generator form for Markovian CPTP evolution. It is central to POS-CP and to OGGE’s claim that open-system dynamics belongs in the theory’s syntax. First cited in §26.

Lisman, J. E., & Jensen, O. (2013). The theta–gamma neural code. *Neuron*, 77, 1002–1016.

Relevance. The canonical hippocampal theta–gamma cross-frequency code, nesting at roughly 1:5–1:8 — a non-dyadic integer ratio, showing the dyadic ($\lambda = 2$) BFFT ladder is a minimal idealization, not a prediction of binary nesting. First cited in §51.

Lorenz, R. D., Lunine, J. I., Withers, P. G., and McKay, C. P. (2001). Titan, Mars and Earth: entropy production by latitudinal heat transport. *Geophysical Research Letters*, 28, 415.

Relevance. The comparative-planetology success of the two-box MEP model worked out in the relevant section of that paper. First cited in §23.

Maes, C. and Netočný, K. (2007). Minimum entropy production principle from a dynamical fluctuation law. *Journal of Mathematical Physics*, 48, 053306.

Relevance. The modern derivation and delimitation: MinEP as the leading order of the occupation-time rate functional around detailed balance, with the frenetic contribution as the obstruction at higher order. First cited in §22.

Maes, C. and Netočný, K. (2013). Minimum entropy production principle. *Scholarpedia*, 8(7), 9664.

Relevance. Authoritative encyclopedic summary by the authors of the fluctuation-theoretic derivation; source for the consolidated validity conditions adopted in this Part. First cited in §22.

Maldacena, J. M. (1998). The large N limit of superconformal field theories and supergravity. *Advances in Theoretical and Mathematical Physics*, 2, 231–252.

Relevance. Maldacena’s AdS/CFT conjecture is the modern archetype for bulk–boundary duality and motivates the idea that geometry can encode quantum gauge dynamics. We cite it as a historical anchor for the GGEP/OGGE posture: bulk geometry is not merely a classical background, but a structure from which boundary quantum physics and observables can be exported. First cited in §1.

Mallat, S. (2009). *A Wavelet Tour of Signal Processing* (3rd ed.). Academic Press.

Relevance. Wavelet and multiscale decomposition theory provides the mathematical background for the BFFT dyadic protocol, which organizes multiscale memory and replay in the Qi-HAL architecture. First cited in §51.

Mandelbrot, B. B. (1982). *The Fractal Geometry of Nature*. W. H. Freeman.

Relevance. Mandelbrot established fractal geometry as a mathematical language for scale-invariant natural structures; the Qi-CON pillar derives fractal attractor supports and anomalous dimensions as consequences of the open-system generator rather than as phenomenological observations. First cited in Part V.

Maturana, H. R., & Varela, F. J. (1980). *Autopoiesis and Cognition: The Realization of the Living*. Reidel.

Relevance. The autopoiesis concept—a self-producing network that maintains its own boundary—is the biological precursor of the SeOS persistence class in the present framework. First cited in §56.

Maupertuis, P.-L. M. de (1744). Accord de différentes lois de la nature qui avoient jusqu’ici paru incompatibles. *Mémoires de l’Académie Royale des Sciences de Paris*, 417–426.

Relevance. The first mechanical least-action statement — mass times speed times distance minimized in all natural processes — read by its author as evidence of design. Cited as the historical hinge between Fermat’s optics and Euler’s mathematics, and as the origin of the overreading habit that has shadowed every extremal principle since. First cited in §21.

Maxwell, J. C. (1865). *A dynamical theory of the electromagnetic field*. Philosophical Transactions of the Royal Society of London, 155, 459–512.

Relevance. The founding paper of classical electromagnetism and the first gauge theory in physics; the conserved current $\partial_\mu J_e^\mu = 0$ is the structural ancestor of the four-current lock of Qi-QUA. First cited in Part III.

Maynard Smith, J., & Szathmáry, E. (1995). *The Major Transitions in Evolution*. W. H. Freeman.

Relevance. The major-transitions framework organizes evolutionary history by successive information-processing innovations; the pagoda hierarchy in the present paper can be read as a physics-grounded restatement of this program. First cited in §66.1.

McCulloch, W. S., & Pitts, W. (1943). *A logical calculus of the ideas immanent in nervous activity*. Bulletin of Mathematical Biophysics, 5, 115–133.

Relevance. Founding paper of theoretical neuroscience; first explicit threshold-logic model of distributed memory in a biological system. Structural ancestor of Qi-HAL. First cited in Part VI.

Metzler, R., & Klafter, J. (2000). The random walk’s guide to anomalous diffusion. *Physics Reports*, 339, 1–77.

Relevance. Foundational paper on anomalous diffusion and Lévy walks; the GDOS derivation of bacterial chemotaxis predicts the biased Lévy walk as a geometric consequence. First cited in §40.

Mitchell, P. (1961). *Coupling of phosphorylation to electron and hydrogen transfer by a chemiosmotic type of mechanism*. Nature, 191, 144–148.

Relevance. The chemiosmotic hypothesis: biological ATP synthesis is driven by a transmembrane proton current. The empirical bioenergetic side of the Qi-QUA Arizaka equation at biological temperatures. First cited in Part III.

Nakagaki, T., Yamada, H., & Tóth, A. (2000). Maze-solving by an amoeboid organism. *Nature*, 407, 470.

- Relevance.** Demonstrates that even a unicellular organism can solve network-optimization problems, providing a low-floor example of adaptive open-system behavior that the GDOS framework must accommodate below the neural level. First cited in §54.
- Neal, R. M. and Hinton, G. E.** (1998). A view of the EM algorithm that justifies incremental, sparse, and other variants. In M. I. Jordan (Ed.), *Learning in Graphical Models*, 355–368. Kluwer, Dordrecht.
- Relevance.** Shows expectation–maximization is coordinate descent on one free energy functional; the template for the FEP’s unification of inference and learning under a single objective. First cited in §24.
- Nicolis, G., & Prigogine, I.** (1977). *Self-Organization in Nonequilibrium Systems*. Wiley.
- Relevance.** A classic monograph on dissipative structures and nonequilibrium self-organization. It provides historical and mathematical grounding for the idea that open systems generically develop attractors and phases under driving—a key theme of GDOS. First cited in §1.1.
- Noether, E.** (1918). Invariante Variationsprobleme. *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse*, 235–257.
- Relevance.** Noether’s theorem is the clearest statement that symmetry is not optional: it has measurable conservation consequences. This motivates POS-GI as the first admissibility rule and supports the constitutional emphasis on symmetry. First cited in Part III.
- O’Keefe, J., & Dostrovsky, J.** (1971). The hippocampus as a spatial map: Preliminary evidence from unit activity in the freely-moving rat. *Brain Research*, 34, 171–175.
- Relevance.** The discovery of hippocampal place cells is a foundational observation for allocentric spatial representation; in the Qi-HAL framework, place-like states instantiate the Inner Hologram and the NHT phase-coding architecture. First cited in §64.5.
- O’Keefe, J., & Dostrovsky, J.** (1971). *The hippocampus as a spatial map: preliminary evidence from unit activity in the freely-moving rat*. *Brain Research*, 34, 171–175.
- Relevance.** The discovery of place cells in the rat hippocampus; first explicit demonstration of distributed spatial-memory representation in a population of biological neurons. Empirical anchor of the Qi-HAL program. First cited in Part VI.
- Onsager, L.** (1931). Reciprocal relations in irreversible processes. I. *Physical Review*, 37, 405–426.
- Relevance.** Onsager’s reciprocity and variational viewpoint provide a foundational extremal structure for near-equilibrium nonequilibrium thermodynamics. We cite it as historical and mathematical background for POS-2 (UEP), supporting the statement that open-system dissipation can be generated by a variational principle (minimum dissipation / gradient flow). First cited in §17.
- Onsager, L., & Machlup, S.** (1953). Fluctuations and irreversible processes. *Physical Review*, 91, 1505–1512.
- Relevance.** Introduces an action-like functional for stochastic trajectories near equilibrium, providing a bridge between variational principles and stochastic thermodynamics. It supports

- POS-2 (UEP) in the open-system regime by showing how “least action” persists as a most-probable-path principle for noisy dynamics. First cited in §21.
- O’Regan, J. K., & Noë, A.** (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, 24, 939–1031.
Relevance. Sensorimotor-contingency theory: perceptual content is constituted by lawful action-dependent changes in sensory input — the cognitive-science form of the action-conditioned observability used here. First cited in §64.4.
- Ozawa, H., Ohmura, A., Lorenz, R. D., and Pujol, T.** (2003). The second law of thermodynamics and the global climate system: a review of the maximum entropy production principle. *Reviews of Geophysics*, 41, 1018.
Relevance. The canonical geophysical review of the climate MEPP, its planetary tests, and its turbulence applications. First cited in §23.
- Paltridge, G. W.** (1975). *Global dynamics and climate — a system of minimum entropy exchange*. Quarterly Journal of the Royal Meteorological Society, 101, 475–484.
Relevance. Founding paper of the maximum-entropy-production tradition (MEPP); applied to the atmospheric system but generalized to all far-from-equilibrium open systems. One of the four gauges that Qi-EP unifies. First cited in §23.
- Paltridge, G. W.** (1978). The steady-state format of global climate. *Quarterly Journal of the Royal Meteorological Society*, 104, 927–945.
Relevance. Restates the 1975 extremum as maximum entropy production of the internal transport, the canonical form of the climate MEPP. First cited in §23.
- Parr, T., Pezzulo, G., and Friston, K.** (2022). *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior*. MIT Press, Cambridge, MA.
Relevance. The consolidated textbook treatment of active inference as a process theory and modeling framework, the stratum of the program that stands independently of the physics-phase claims. First cited in §24.
- Polettini, M.** (2013). Fact-checking Ziegler’s maximum entropy production principle beyond the linear regime and towards steady states. *Entropy*, 15, 2570–2584.
Relevance. The constitutive-versus-phenomenological dichotomy that separates what the flux-side and force-side readings each select, with counterexamples beyond linear response. The external statement of the boundary that Theorem 30.1 derives. First cited in §23.
- Prigogine, I. and Lefever, R.** (1968). Symmetry Breaking Instabilities in Dissipative Systems. II. *Journal of Chemical Physics*, 48, 1695–1700.
Relevance. The Brusselator: explicit demonstration that beyond a finite distance from equilibrium the stable regime is not the branch of least entropy production. First cited in §22.
- Prigogine, I.** (1945). Modération et transformations irréversibles des systèmes ouverts. *Bulletin de la Classe des Sciences, Académie Royale de Belgique*, 31, 600–606.

Relevance. The priority document: first statement and proof of the minimum entropy production theorem, including the monotonic decrease of P during relaxation. First cited in §22.

Prigogine, I. (1947). *Étude thermodynamique des phénomènes irréversibles*. Desoer, Liège.

Relevance. The doctoral thesis in book form and the theorem’s canonical early exposition, containing the thermodiffusion example of the relevant section of that paper. First cited in §22.

Prigogine, I. (1967). *Introduction to Thermodynamics of Irreversible Processes*. Interscience.

Relevance. Prigogine’s exposition (including the minimum entropy production theorem near equilibrium) is a standard historical source for entropy-production variational ideas. We cite it in POS-2 (UEP) to connect the PLA-to-dissipation bridge with classical nonequilibrium thermodynamics and to clarify when entropy-production extremals are mathematically justified. First cited in §1.1.

Prigogine, I. (1978). Time, Structure and Fluctuations (Nobel lecture, December 8, 1977). *Science*, 201, 777–785.

Relevance. Prigogine’s own retrospective, presenting the linear-regime theorem together with the breakdown of its variational structure as the gateway to dissipative structures. First cited in §22.

Raja, V., Valluri, D., Baggs, E., Chemero, A., and Anderson, M. L. (2021). The Markov blanket trick: on the scope of the free energy principle and active inference. *Physics of Life Reviews*, 39, 49–72.

Relevance. Argues the blanket is induced by the modeler’s parcellation rather than discovered in the system, with the Watt governor as central counterexample to universal scope. First cited in §24.

Ramstead, M. J. D., Sakthivadivel, D. A. R., Heins, C., Koudahl, M., Millidge, B., Da Costa, L., Klein, B., and Friston, K. (2023). On Bayesian mechanics: a physics of and by beliefs. *Interface Focus*, 13(3), 20220029.

Relevance. The programmatic statement of Bayesian mechanics, comparing the FEP to a principle of least action on belief spaces and making the deflationary reading official. First cited in §24.

Rao, R. P. N. and Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1), 79–87.

Relevance. The canonical hierarchical predictive-coding model that the Gaussian gradient dynamics of the FEP reproduce; an algorithm that neither requires nor implies the principle. First cited in §24.

Rayleigh, Lord (J. W. Strutt) (1873). Some General Theorems relating to Vibrations. *Proceedings of the London Mathematical Society*, 4, 357–368.

Relevance. Introduces the dissipation function for linearly damped systems; Onsager’s variational function is modeled on it. First cited in §22.

Schnakenberg, J. (1976). Network theory of microscopic and macroscopic behavior of master equation systems. *Reviews of Modern Physics*, 48, 571–585.

Relevance. Defines the standard entropy production functional for Markov jump processes, used in the stochastic formulation of MinEP. First cited in §22.

Schrödinger, E. (1944). *What Is Life? The Physical Aspect of the Living Cell*. Cambridge University Press.

Relevance. Schrödinger’s question “What is life?” frames the central challenge that GDOS with Qi-4 Pillars addresses: how a physical framework can account for the persistence, adaptation, and information-processing capacities of living systems without invoking new fundamental laws beyond those governing the open-system boundary. First cited in §1.1.

Seifert, U. (2012). Stochastic thermodynamics, fluctuation theorems and molecular machines. *Reports on Progress in Physics*, 75, 126001.

Relevance. Stochastic thermodynamics extends thermodynamic concepts to small fluctuating open systems and provides precise expressions for entropy production and nonequilibrium work. We cite it for the appendix toy models that connect GDOS-style open dynamics to measurable entropy production and for motivating POS-2’s open-system extremal forms. First cited in §17.

Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27, 379–423, 623–656.

Relevance. Shannon’s information theory provides a quantitative language for “entropy” and information flow. Because GG-Constitution aims to export a single operational grammar across sciences, information-theoretic quantities become natural candidates for inclusion in the open-system ledger (GDOS) implied by OGGEF. First cited in §1.1.

Simon, H. A. (1962). The architecture of complexity. *Proceedings of the American Philosophical Society*, 106, 467–482.

Relevance. Simon’s “nearly decomposable” hierarchical architecture is a conceptual ancestor of the pagoda hierarchy; the present paper gives it a geometric runtime basis through the story-lift theorem and the DGI shinbashira. First cited in §72.1.

Singh, M. S., and O’Neill, M. E. (2022). The climate system and the second law of thermodynamics. *Reviews of Modern Physics*, 94, 015001.

Relevance. The modern consensus assessment: theoretical and modeling issues preclude broad acceptance of MEPP in climate science. First cited in §23.

Spohn, H. (1978). Entropy production for quantum dynamical semigroups. *Journal of Mathematical Physics*, 19, 1227–1230.

Relevance. The entropy-production theorem: $-dD(\rho_t \parallel \rho_*)/dt \geq 0$ for any CPTP semigroup with stationary state ρ_* . Monotonicity follows from stationarity alone; the free-energy reading needs the detailed balance of Kossakowski et al. (1977) in addition. First cited in §8.

Stinespring, W. F. (1955). Positive functions on C^* -algebras. *Proceedings of the American Mathematical Society*, 6, 211–216.

- Relevance.** Stinespring’s dilation theorem is the mathematical backbone of quantum channels: CP maps arise from unitary evolution on an enlarged space. We cite it to justify how OGGEp can treat open-system boundary dynamics as fundamental while remaining consistent with microscopic unitarity in the bulk. First cited in §9.
- Susskind, L.** (1995). *The world as a hologram*. *Journal of Mathematical Physics*, 36, 6377–6396.
- Relevance.** Sharpened the holographic principle of ‘t Hooft (1993) into the covariant entropy bound. Structural input to the OH+IH=EIC framework. First cited in Part VI.
- ’t Hooft, G.** (1993). *Dimensional reduction in quantum gravity*. In “Salamfestschrift” (eds. Ali, Ellis, Randjbar-Daemi), World Scientific, pp. 284–296. arXiv:gr-qc/9310026.
- Relevance.** Founding statement of the holographic principle: the information content of a region of space is bounded by the area of its boundary. Structural ancestor of OH+IH=EIC. First cited in §1.
- Tononi, G.** (2004). An information integration theory of consciousness. *BMC Neuroscience*, 5, 42.
- Relevance.** IIT’s integrated information measure Φ is compared with GDOS’s Qi-HAL universality theorem in the consciousness chapters. First cited in Part VI.
- Touchette, H.** (2009). The large deviation approach to statistical mechanics. *Physics Reports*, 478, 1–69.
- Relevance.** A modern review of large-deviation theory, clarifying how action-like functionals control typical and rare trajectories. It supports POS-UEPby providing a general mathematical framework for extremal principles in open systems. First cited in §7.
- Turner-Evans, D., et al.** (2020). The neuroanatomical ultrastructure and function of a biological ring attractor. *Neuron*, 108, 145–163.
- Relevance.** Reviews neural ring attractor circuits in *Drosophila*, providing a direct biological realization of the persistent phase-coded memory and evidence-dependent switching described by GA and the DGI–GSSB switch in the Qi-HAL framework. First cited in Part VI.
- van Kampen, N. G.** (2007). *Stochastic Processes in Physics and Chemistry* (3rd ed.). Elsevier.
- Relevance.** Develops stochastic-process tools used for mesoscopic physics and chemical kinetics. It provides background for interpreting GDOS as the universal open-system bridge between microscopic laws and coarse-grained dynamics. First cited in §4.1.
- Vicsek, T., Czirók, A., Ben-Jacob, E., Cohen, I., & Shochet, O.** (1995). Novel type of phase transition in a system of self-driven particles. *Physical Review Letters*, 75, 1226–1229.
- Relevance.** Introduces a minimal model of collective self-driven motion that has become the standard reference for emergent flocking; in the GDOS framework it instantiates the collective-level DGI–GSSB–GA pipeline applied to distributed social or biological systems. First cited in §O.
- von Holst, E., & Mittelstaedt, H.** (1950). Das Reafferenzprinzip. *Naturwissenschaften*, 37, 464–476.

- Relevance.** The reafference principle: the nervous system separates self-produced (reafferent) from externally produced sensory change via an efference copy — the mechanistic basis for why action is constitutive of internal-model formation. First cited in §75.3.
- von Neumann, J.** (1932). *Mathematische Grundlagen der Quantenmechanik*. Springer.
- Relevance.** von Neumann formalized quantum theory axiomatically but introduced measurement as an additional process distinct from unitary evolution. This motivates GG-Constitution’s insistence that measurement be defined as a constitutional axiom/postulate rather than left as a dualistic rule. First cited in §1.2.
- Wang, J. C.** (1979). Helical repeat of DNA in solution. *Proceedings of the National Academy of Sciences of the USA*, 76, 200–203.
- Relevance.** Establishes the helical repeat of DNA in solution (~ 10.5 bp/turn), a structural integer that illustrates how geometric periodicity constrains the molecular-scale SIP and the topology of the structural-material current J_B . First cited in §17.
- Watson, J. D., & Crick, F. H. C.** (1953). Molecular structure of nucleic acids: A structure for deoxyribose nucleic acid. *Nature*, 171, 737–738.
- Relevance.** The double-helical structure of DNA is the foundational molecular template whose geometric and informational properties motivate the chemistry-to-life transition in the GDOS pagoda hierarchy. First cited in §17.
- Wilson, K. G.** (1971). Renormalization group and critical phenomena. I. *Physical Review B*, 4, 3174–3183.
- Relevance.** Wilson’s RG is the prototype of scale-covariant thinking and motivates POS-RG. Because GG-Constitution aspires to unify across scales and domains, it demands that coarse-graining and operational measurement commute in a controlled way. First cited in Part V.
- Wilson, K. G., & Kogut, J.** (1974). The renormalization group and the ϵ -expansion. *Physics Reports*, 12, 75–199.
- Relevance.** A standard reference for the modern RG viewpoint: coarse-graining as a flow on theory space. POS-5 builds on this picture but adds the constitutional demand that RG must commute with the operational dictionary (OGGEP), enabling a geometric RG coordinate in GG-6. First cited in Part V.
- Yoshida, M., Muneyuki, E., & Hisabori, T.** (2001). ATP synthase—a marvellous rotary engine of the cell. *Nature Reviews Molecular Cell Biology*, 2, 669–677.
- Relevance.** Authoritative review of ATP synthase; the rotary mechanism and energy budget are the empirical data against which the Arisaka Equation is verified. First cited in §17.
- Ziegler, H.** (1961). Zwei Extremalprinzipien der irreversiblen Thermodynamik. *Ingenieur-Archiv*, 30, 410–416.
- Relevance.** Companion statement of the two extremum principles (force space and flux space) whose flux-space member is the maximum entropy production form used throughout this Part. First cited in §23.

Ziegler, H. (1983). *An Introduction to Thermomechanics*, 2nd ed. North-Holland, Amsterdam.

Relevance. Textbook consolidation of the orthogonality principle and its plasticity applications; the canonical reference for the orthogonality condition as this Part states it. First cited in §23.

Ziman, J. M. (1956). The general variational principle of transport theory. *Canadian Journal of Physics*, 34, 1256–1273.

Relevance. States and proves the one uncontested MaxEP theorem: the stationary solution of the linearized Boltzmann equation maximizes entropy production under the self-consistency constraint. First cited in §23.

Zurek, W. H. (2003). Decoherence, einselection, and the quantum origins of the classical. *Reviews of Modern Physics*, 75, 715–775.

Relevance. Decoherence provides a mechanism for effective classicality in open systems and illustrates why openness and environment-induced selection are essential to operational measurement. This motivates OGGEp and the elevation of open dynamics to the constitutional layer. First cited in §1.2.