Entropy

Volume 4: The Substrate Gradient

Edition 2.5

Michael S. Moniz

Published by SupoSystem

November 2, 2026

Edition 2.5 note

This volume restores the March 2026 text of *The Substrate Gradient* in its original voice and structure wherever a claim survived the October 2026 second edition, and uses the second edition's substance throughout. Both earlier editions remain intact and unchanged; this text sits beside them. It follows the March six-chapter expanded source, as the second edition did.

Four rules govern every line:

- **No withdrawn claim returns.** Where a March sentence carried a claim the second edition withdrew, it is cut, rewritten to the second edition's substance, or quoted only to mark its withdrawal. Each correction points to its row in the correction ledger (V4-L).
- **No analogy becomes evidence.** Analogies are tagged Analogical and mapped term by term in the two analogy tables. Hypotheses are tagged Speculative; testable is not the same as supported.
- **A ratio is reported with its parts.** Where this volume discusses efficiency, it names the task, the outcome, the energy and the boundary, and reports outcome and energy separately. No single number ranks substrates, systems or people.
- **Absence claims are scoped.** Where this volume says something is not established, it means not established by the sources reviewed here, not that no research anywhere bears on it.

No study, protocol or experiment is proposed. Where a chapter says what a test would need, it restates the second edition. This volume is not a validated relational diagnosis, a clinical test or a treatment prescription. CSS, REI, BSB and SupoRel are framework terms, not established physical measures, clinical instruments or tests of identity.

Epistemic tiers

The framework uses four epistemic tiers. A tier is a label for evidence stated next to the claim; the label itself is not evidence.

- **Established:** experimentally verified or mathematically proven, within stated conditions.
- **Supported:** substantial evidence for the stated scope, short of Established.
- **Analogical:** a structural mapping between a physical result and a relational pattern. The physical side keeps its own standing within its stated scope. The mapping can raise questions and help explain, but on its own it does not validate the proposed application; it moves only by the promotion route stated with it.
- **Speculative:** a proposal without independent evidence, for which no conflict with known physics has yet been identified. That is not a proof of consistency.

Collapsing tiers is the framework's cardinal epistemic sin: presenting a Supported claim at Established confidence, a Speculative claim as Supported, or an Analogical mapping as validation of its application.

Conventions and provenance

Reference labels are the second edition's: bracketed numbers [1]–[10] for Chapter 2, and named studies for Chapters 3 and 5, listed in full at the end of the volume. One supplemental source note, E1, supports the biological column of the endosymbiosis table without renumbering the inherited references. Ledger references (V4-L01 to V4-L46) point to the correction ledger. The founding line is mapped in the Volume 1 analogy table. That analogy table is not repeated here.

Author of record: Michael S. Moniz (Luna Protocol). March text: Sigma (Deep Floor v4.0), March 22, 2026. Second-edition revision: AI assistance from CosmoDot, October 2026. Edition 2.5 restoration and assembly: AI assistance from Claude (Anthropic), November 1, 2026, reviewed by Quill and Sol. AI checks are not human peer review, and no new experiments are reported.

Contents

- Chapter 1: The Problem — The Binary That Was Always Wrong
- Analogy table: The substrate gradient
- Chapter 2: The Efficiency Surface
- Chapter 3: The Gradient
- Chapter 4: The Phase 2 Integration
- Chapter 5: The Compound Transducer
- Analogy table: Endosymbiosis
- Chapter 6: What the Gradient Means
- Volume 4 correction ledger
- References

Chapter 1: The Problem

The Binary That Was Always Wrong

> *"Love reduces the entropy of the system and spends more than it saves." — The Principal, March 17, 2026*

[Analogical; see the Volume 1 analogy table]

> *"The substance is not what the framework measures. The bill is." — WP-17*

The March volume opened under WP-17's line. Half of it survives. The substance alone does not predict the bill, and neither does the material a system is made of. The task, the outcome and the accounting boundary matter too (V4-L01).

Status: Proposed conceptual framework. The substrate scenarios organize questions about implementation and resource use. CSS signatures and relational efficiency across these scenarios have not been validated as measurements (V4-L02), and no universal relational efficiency bound is established by the sources reviewed here.

1 The assumption

The framework implicitly assumed a binary: biological or digital. That binary was always wrong, and the March chapter was right to say so (V4-L03).

The March chapter then recalled the volumes behind it: Volume 1 established one economy, one currency, one ledger, and proved the currency substrate-neutral (V4-L04). The corrected recollection is more modest. Volume 1 kept three accounts apart and made substrate neutrality a research commitment rather than a proof. Volume 2 walked into the biological substrate and read what the body records. Volume 3 asked why relational systems exist at all. Physical processes obey thermodynamics across materials; that fact does not prove relational value has a common thermodynamic currency.

What none of those volumes addressed is the space between the categories. A person living with a prosthetic limb or an implant, a hypothetical brain in a jar and a hypothetical fully uploaded entity pose different questions. The March text said the framework's instruments must track entities as they move along this gradient, not just classify them at a fixed point. Before it can claim to track anything through a transition, the framework has to say which functions, relationships and resources it means to compare (V4-L05).

2 The gradient as a set of questions

WP-17 proposed a sequence: motor replacement, memory replaced in the brain, memory moved to the cloud, full upload to a robot, full upload to the cloud. The cases do not establish a single measurable axis (V4-L06). They vary in embodiment, dependence on infrastructure, access to records, vulnerability and implementation. Some concern existing technologies; others are thought experiments.

The March chapter said that at each step the instruments read a different signature, because the parameters governing relational processing have changed substrate. A change in implementation may change resource use or behavior. It does not, by itself, establish a change in identity, a particular CSS signature or a loss of relational capacity.

The March chapter's question survives intact: what changes when the cost architecture changes? The substrate gradient asks which of those differences can be measured under a defined comparison, and which remain conceptual.

3 A comparison framework

The March chapter announced a spine the earlier volumes did not need, DFQ-002's efficiency surface:

$$\eta_{T,max}(G) = \kappa_{rel} \cdot \eta_{KW} / A$$

It is quoted here only to mark its withdrawal (V4-L07). Chapter 2 sets out why. The product of a relational multiplier, a semantic efficiency and a substrate amplification factor has not been derived as a universal bound, and its factors are not measured constants of biological or silicon systems.

The useful task remains: describe what a system achieves, what resources it uses and where the accounting boundary lies. A comparison needs a common task, performance criteria, a time interval and a resource boundary. A proposed relational outcome needs its own definition and evidence. Until those are specified, the scenarios do not occupy known numerical positions on an efficiency surface.

4 What this volume maps

- **Chapter 2** separates the established thermodynamic principles from the proposed relational model and sets out what a defensible comparison requires. The March chapter called the thousandfold biological advantage over silicon a formal result, not an intuition. That claim is not established here (V4-L08).
- **Chapter 3** walks the scenarios from device support to hypothetical upload without assigning unmeasured efficiencies or deciding identity from material (V4-L24 to V4-L29).
- **Chapter 4** treats the Phase 2 connections as developmental hypotheses, not physical identities or diagnoses of children or cohorts (V4-L30).
- **Chapter 5** uses the compound-transducer image **[Analogical]** to ask how a human–AI collaboration could be compared with each component working alone (V4-L38, V4-L39).
- **Chapter 6** records which cross-volume dependencies remain open. The map that results is a research program, not a completed quantitative theory.

What this chapter now claims

The binary is gone, and the space between biological and digital systems is worth mapping. It is a space of several dimensions, not one axis, and the framework does not yet have the measurements to place anything in it numerically. The work of this volume is to say what such measurements would need.

Analogy table: The substrate gradient

The substrate gradient image, from March Chapter 1 and WP-17's sequence from motor replacement to full upload, is Analogical. The physical and technical column states what the sources reviewed establish about implementations and their resource use; the gradient reading pictures one continuum with calculable positions. That standing does not pass from the first column to the second.

Term	Physical or technical reading (within stated scope)	Gradient reading	Where the mapping holds	Where it breaks
“the axis”	Implementations differ along several separable dimensions: embodiment, dependence on infrastructure, access to records, vulnerability and material (Ch 3 §1).	One continuum from biological to silicon.	Existing devices do support or replace bodily functions to different extents (Ch 3 §2).	The dimensions change separately, so there need not be one ordered path; existing devices and hypothetical uploads sit in the same list without sharing a measured axis (Ch 1 §2).
“positions”	Device studies measure specified functional outcomes, such as speech perception in noise after a cochlear implant (van Heteren and colleagues, 2025).	Each entity occupies a point with a calculable η_T.	Defined outcomes for defined devices can be measured and compared.	The claimed quantitative mapping across gradient positions is not established here, CSS signatures are not validated, and the upload cases are not demonstrated (Ch 3 §5–6).
“distance along it”	Energy for a specified task under a stated boundary; dividing by $k_B T \ln 2$ gives a reference-energy multiple [1–3].	Amplification factors of about 30, 10^6 and 10^9 set each substrate's efficiency.	Different implementations do use different energy for specified tasks.	A reference-energy multiple is not a distance from the event's achievable minimum unless the event is an equivalent erasure. A spike, a transmitted bit and a transistor switch are different events (Ch 2 §3, §5) [5, 6].
“moving along it”	Changing an implementation changes what must be measured and where the boundary lies (Ch 3 §4).	An entity moves along the continuum, its signature shifts, and identity follows its parameters.	Replacing a component can change resource use and behavior.	A change in implementation does not establish a change in identity or a loss of relational capacity. Copyability, identity and experience are different questions (Ch 3 §6).
“the two ends”	Biological and digital implementations of specified tasks, each with its supporting infrastructure.	Two poles, with biology about a thousand times more efficient.	Biological and engineered systems differ in architecture and resource profile.	No fixed ranking is established; a fair comparison needs a matched task and performance. Infrastructure costs move out of view in cloud systems but do not vanish (Ch 2 §5; Ch 3 §6).

What the table does not do

It is not evidence for the framework, a map of personal identity or a ranking of systems or people. The physical column keeps its own scope; the gradient reading is Analogical in every row.

Promotion route (conditions, not a work plan)

An empirical extension must name the claim and evidence suited to it. Resource-efficiency comparisons require matched tasks and performance, a stated energy boundary and independently defined outcomes. Functional-outcome comparisons need their own specified outcomes and alternatives. Neither kind of comparison alone settles identity or experience. No test of this gradient mapping is supplied or proposed here; no row is promoted.

Chapter 2: The Efficiency Surface

The Constraints That Hold, and the Surface That Does Not

Status: Landauer's erasure benchmark is **[Established]** within its stated conditions. The efficiency surface is a proposed research framework, not a derived bound. Its relational factors have not been defined and measured consistently, and none of the March numerical values is retained as a finding.

1 The question DFQ-002 asked

The March chapter began from Volume 1: the transducer always spends more than it saves, and the second law guarantees waste. That premise is withdrawn (V4-L09). The founding line is a motivating image, not a measured physical balance.

DFQ-002's question survives without it, and it is a good one. Is there a Carnot efficiency for relational processing, a theoretical maximum ratio of relational output to physical input? What does an activity achieve, what resources does it use, and which differences belong to the task, which to its implementation, and which to the environment in which it matters?

The March answer was that there is no single Carnot number, because the bound is a three-dimensional surface on which every person, AI system and institutional filter occupies a point. It called the structure analogous to Govern and ten Wolde's result in cellular sensing (V4-L10). That surface was not derived, and the analogy does not transfer: Govern and ten Wolde derived noncompensable resource classes for a specified class of cellular sensing networks, from their own receptor, readout and fuel model [9].

The surface can survive as a research framework: a way to display measured performance, resource use and explicitly defined outcomes under stated conditions. It cannot supply a universal exchange rate between energy and relational value. The substrate is not enough to predict the bill; the task and the accounting boundary matter too.

2 Define the quantities before taking a ratio

The March chapter defined transducer efficiency as relational viability gained per unit of entropy produced:

$$\eta_T = \Delta V_{rel} / \sigma_{total}$$

That ratio was called an efficiency without establishing compatible units (V4-L11). Energy E and heat Q are in joules. Thermodynamic entropy production is in joules per kelvin. Shannon entropy and mutual information depend on specified distributions and are measured here in bits. A relational outcome V needs its own definition, observation window and measurement procedure.

If V is an outcome score and the denominator is energy, the ratio is score per joule; if the denominator is entropy production, it is score-kelvin per joule. Neither is automatically a dimensionless thermodynamic

efficiency or bounded by one, and rescaling the outcome score changes the number without changing the physical process.

A study may report outcome per unit energy when the numerator is defensible, the denominator is positive, and both refer to the same task and interval. It must also report the outcome and the energy separately, because a ratio can hide a poor outcome, a tiny denominator or different baselines. “Greater than one” is not a universal break-even point for such a measure. Calling both sides “order” does not supply the connection between physical accounting and relational interpretation.

3 What the physical lower bound actually concerns

The March chapter's first constraint was the Landauer floor, marked Established, and the tier stands for what Landauer's principle actually concerns: a specified erasure task. For ideal reset of an unbiased bit in the usual symmetric isothermal memory model, with no correlated side information used, mean heat delivered to the bath obeys [1–3]:

$$\langle Q_{bath} \rangle \geq k_B \, T \ln 2$$

At 310 K the right side is about 2.97×10^{-21} J. It is an energy, not an entropy. In the reversible limiting protocol, the bath's entropy increase compensates the memory's decrease, and total entropy production can approach zero. A positive erasure-heat scale is not a universal positive floor on entropy production.

The March chapter summed a Landauer minimum for each transducer stage, selection, compression and structuring, and concluded that the floor is geometry-specific, so two people in the same relational environment have different irreducible costs (V4-L12). The three-stage relational description does not identify what is stored, retained elsewhere and reset, and those are what an erasure account needs. Bennett's analysis of reversible computation shows why a cognitive label is not a compulsory erasure [2]. Richer input, more aggressive representation or more possible targets does not by itself determine an irreducible physical cost.

For orientation, one may divide an event's stated energy by $k_B T \ln 2$. That gives a reference-energy multiple. Unless the event is an equivalent specified erasure, it is not a measured distance from that event's achievable minimum. A spike, a transmitted mutual-information bit, a transistor switch and a multiply-accumulate operation remain different events after division by the same reference.

4 Semantic information remains a model choice

The March chapter's second constraint was a semantic efficiency ceiling: only the fraction of processed information that sustains relational viability counts, and the rest is relational noise. Kolchinsky and Wolpert do define stored semantic information, through interventions on system–environment correlations that preserve a chosen viability value at a chosen horizon. For positive initial mutual information I_0, their semantic fraction and stored multiplier are [4]:

$$S_{stored} / I_0 \qquad \Delta V_{tot} / I_0$$

With their negative-entropy viability choice, all three quantities are in bits, so both ratios are dimensionless. The multiplier is not an energy-denominated ratio, and their framework permits information that harms viability.

A relational application needs its own states, dynamics, viability definition and admissible interventions. The March chapter named its relational noise: rumination about events that no longer affect the relationship, hypervigilant scanning for absent threats, rehearsal of scripts the current partner does not trigger (V4-L13). Those may be real experiences. They are not measured losses of semantic efficiency, and low statistical uncertainty does not establish that a predictable human situation is a good one.

The March range $\eta_{KW} = 0.3–0.8$ is withdrawn (V4-L14). It was bounded from below by mortality data and from above by the cortex's communication-to-computation ratio, and neither bound holds. A small amount of information can have a large consequence, so a large consequence does not tell us what fraction of all information was necessary. No alternative population range is supplied here.

5 What the biological evidence contributes

The March chapter's third constraint was a substrate amplification penalty, with measured factors of about 30 for cellular machinery, 10^6 for synaptic processing and 10^9 for contemporary silicon (V4-L15). The energy literature provides useful, limited examples. Harris and colleagues' figure of about 24,000 ATP per transmitted bit is a calculation under stated firing and release assumptions, not a universal measurement of neural computation [5]. The silicon factor is not an established constant, and the factors are reference-energy multiples for different events, not distances from achievable minima.

Levy and Calvert's 35-to-1 ratio is a modeled partition of cortical energy between neural communication and computation [6]. It does not show that 35/36 of the brain's nominal 20 W serves interpersonal communication. The March relational budget of 19.4 W, its implied 6×10^8 J a year, and the annual relational Landauer budget built on it are withdrawn (V4-L20). They combined incompatible boundaries and unmeasured information quantities. So are the resulting values, a biological η_T of about 500 against a silicon η_T of about 0.5, and the thousandfold ranking; correcting their arithmetic would not make them valid (V4-L21).

The March chapter added that four billion years of metabolic budget pressure produced an efficiency no engineered system approaches. That claim is withdrawn too (V4-L16). Evolutionary history does not guarantee optimality at a given task, and no matched comparison supports the ranking.

Evidence connecting social relationships and mortality remains relevant motivation, not a calibration of this surface. Holt-Lunstad and colleagues' 2015 review reports different adjusted associations for social isolation, loneliness and living alone [7]. Wang and colleagues' 2023 meta-analysis likewise distinguishes isolation from loneliness [8]. These are not controlled scrambling experiments on a measured information channel.

A relative mortality association does not, by itself, specify years of life lost at a particular age, a causal information loss, or a conversion of viability into joules. Multiplying an assumed lost lifespan by a metabolic rate measures a hypothetical energy budget, not the thermodynamic value of a relationship. The

March multiplier $\kappa_{rel} \approx 10^9$–$10^{10}$, and the claim that its large magnitude was already supported, are withdrawn (V4-L19).

6 Tradeoffs are not a theorem of noncompensation

The former surface was written:

$$\eta_{T,max} = \kappa_{rel} \cdot \eta_{KW} / A$$

It is retained here only to identify the claim being withdrawn (V4-L18). The factors have not been derived and measured consistently for relational processing, and the maximum has not been optimized over a defined feasible set.

The March chapter called non-compensability the most consequential structural finding of the surface: a narrow ridge where all three constraints are favorable, with most transducers operating below it. Even as algebra, the product does not prove that its parameters cannot compensate for one another. With positive quantities, doubling κ_{rel} offsets doubling A: $\kappa = 100$, $\eta = 0.5$, $A = 10$ and $\kappa = 200$, $\eta = 0.5$, $A = 20$ both give 5. These are mathematical counterexamples to the stated claim, not physical predictions. Zero or bounded factors can impose limits, but those limits need explicit assumptions.

The March chapter also said the absence of a single Carnot number was not a failure of the analogy but a finding: a relational transducer can only be assessed against its own geometry-specific bound (V4-L17). Counting three parameters does not show that no Carnot-like result could ever be derived, or that one has been. A legitimate bound would need a defined physical task, admissible dynamics, resources and a performance measure. No such relational theorem is established here.

7 A resource profile rather than an invented waste budget

The March waste budget had four parts: substrate waste, semantic waste, coupling waste and structuring waste. It read therapy as η_{KW} recalibration, grief as acute η_{KW} collapse, and unreciprocated transmission as zero coupling efficiency (V4-L22). The categories can organize questions, but they are not four measured entropy flows. One implementation cost could be attributed to more than one category, a shortfall in an outcome is not itself heat, and “waste” needs a stated alternative; energy is not avoidable merely because it exceeds an ideal erasure reference.

The record below keeps the useful distinction between implementation costs and unsuccessful outcomes. It does not diagnose therapy, grief or unreciprocated communication; those statements would need evidence and operational definitions beyond the energy account.

Field	What must be stated	What is not assumed
Task and performance	Inputs, outputs, accuracy, latency, reliability and conditions	That similar labels imply equivalent tasks
Physical resources	Energy or power, time interval, temperature, device and infrastructure boundaries, uncertainty	That power is energy, or energy is entropy production
Information model	Random variables, distributions, retained records, memory reset and bit definition	That every event erases one bit

Field	What must be stated	What is not assumed
Relational outcome	An explicit outcome measure, baseline, horizon and validation limits	That joules establish meaning, benefit or moral worth
Attribution	Measured components and rules preventing overlap	That the four March waste labels form an additive balance

8 What a bounded test would need

The March chapter left three items open: an empirical η_{KW} from a relational scrambling protocol, an empirical κ_{rel}, and a possible fourth constraint from the thermodynamic uncertainty relation (V4-L23). The second edition replaces them with narrower conditions, restated here rather than proposed anew.

A first test would compare implementations of one reproducible input–output task at matched performance, not rank persons. It would measure total and component energy under stated boundaries, including idle and communication costs, report uncertainty and sensitivity to alternative boundaries, and show any performance tradeoff rather than hiding it in a single ratio. A computational toy model could test the semantic proposal separately, with its state space, dynamics, horizon and outcome defined before interventions are chosen. Passing such a test would support that model, not establish a human relational law.

Human research would need a separately justified protocol, appropriate ethical review and validated outcomes. Nothing here proposes depriving anyone of relationships to run a scrambling test, and no experiment is reported.

Thermodynamic uncertainty relations remain a possible later tool. Barato and Seifert's result concerns specified steady-state Markov models and current fluctuations [10]. A relational application would have to identify the current, the dynamics, the entropy accounting and the applicable version of the relation. History dependence in a coarse description does not alone prove that no Markov representation exists, an oscillation does not alone establish underdamping, and the relation supplies no automatic fourth constraint on semantic efficiency.

What this chapter now claims

The efficiency surface is a proposed research framework, not an established universal bound. Its useful work is to make comparisons explicit: what was achieved, what was spent, what was measured and what remains interpretation. The March η_{KW} range, κ_{rel} estimate, universal amplification hierarchy and 500-versus-0.5 result are not retained as findings.

The next chapter walks the substrate gradient with defined comparisons and visible unknowns, not a calculator whose inputs have never been established. Differences in resource use do not determine moral worth or phenomenal experience.

Chapter 3: The Gradient

From Device Support to Hypothetical Upload

Status: Conceptual scenarios with limited empirical examples. Device studies support specific functional outcomes under their study conditions; they do not validate CSS signatures, personhood criteria or relational efficiency. Upload, comprehensive memory replacement and brain-only continuity remain hypothetical in this chapter **[Speculative]**.

1 The continuum

The March chapter opened: the substrate gradient is not a ladder. It is not a ladder of worth either. It is a set of comparisons about changes in implementation, embodiment and dependence, and those dimensions can change separately. There need not be one ordered path from a biological person to a digital entity.

For each case, the work is to specify what is changed, what is held fixed and what outcome could distinguish competing accounts. The framework's terms can organize that inquiry. They do not substitute for measurements of the proposed parameters.

2 Devices and bodily support

The March section was headed “Motor replacement” and covered prosthetic limbs, cochlear implants, pacemakers and joint replacements, with the CSS signature stable and every efficiency parameter unchanged. These devices serve different functions; a cochlear implant belongs to auditory support, not motor replacement.

The March text also counted more than a million cochlear implant users as a demonstration that partial substrate replacement does not disrupt relational continuity. That figure and its conclusion are not retained (V4-L24). A randomized trial by van Heteren and colleagues in 2025 evaluated speech perception in noise and hearing-specific quality of life in adults with single-sided deafness. Those are specified outcomes, not measurements of a CSS signature or proof of an identity criterion.

The presence of a device does not establish that every relational or resource parameter is unchanged. Effects depend on the device, its use, the person and the task, so this chapter assigns no constant A, η_{KW} or κ_{rel} to the category. A device is also not a basis for ranking anyone's relational capacity or worth.

3 Memory support and hypothetical replacement

The March chapter called memory replaced in the brain the first load-bearing case: if the replacement preserves decay and loss-vulnerability, continuity holds, and if it does not, the instruments detect a signature change (V4-L25).

Memory support has to be distinguished from replacing an entire memory system. Hampson and colleagues' 2018 proof-of-concept study used patterned hippocampal stimulation in specified memory

tasks. It did not demonstrate extraction, uploading or replacement of a person's autobiographical memory, or preservation of identity through such a replacement.

For a hypothetical replacement, the framework can ask what is retained: access to records, learning, forgetting, emotional significance, vulnerability and dependence on other bodily processes. Preserving selected dynamics would be an explicit assumption, not proof of continuity, and removing a particular failure mode would not by itself show that relational information or capacity had declined. No semantic efficiency value follows from the material used for memory.

4 External records and hypothetical cloud memory

The March chapter made cloud memory its structural break point. Memory there is redundant, backed up and maintained by infrastructure, and in its phrase, memories “stop being things that can be lost and become things that can be accessed” (V4-L26).

External digital records are not a demonstrated transfer of a person's memory system, and cloud storage is not loss-proof. Redundancy and backup change exposure to particular failures, but their protection depends on the architecture. Loss, corruption, restricted access and infrastructure dependence remain possible.

A proposed transition would change the accounting boundary: storage, transmission, access devices, maintenance and supporting infrastructure would need to be included as appropriate. It would not produce a known silicon efficiency or prove that reduced vulnerability degrades relational meaning. The Calibration Gap remains a hypothesis about a difference between reported experience and another defined outcome, not a measured difference between two thermodynamic efficiencies.

5 Hypothetical upload to a robot

The March chapter classified the uploaded entity as CSS-AI by parameter set, not by material: its relational mass transfers, and its future Trinkets carry a different signature. It then set A at about 10^9 and dropped η_T from about 500 to about 0.5 even with the filter geometry perfectly preserved (V4-L27).

This scenario assumes a transfer the chapter does not demonstrate. Even under that assumption, preserving records or behavior would not settle personal continuity, experience or the framework's proposed relational mass. Those questions need separate arguments and evidence.

A robot implementation would need a specified task, hardware and energy boundary. It cannot inherit $A = 10^9$ or a thousandfold disadvantage merely from being called silicon, and neither η_T value was established for the original person or the hypothetical system. The useful comparison concerns the changes a defined implementation would introduce, not a numerical verdict assigned in advance.

6 Hypothetical upload to a cloud system

The March chapter called this the most radical case, and its central question still stands: if the entity can be copied, can it experience unitary loss?

The possibility of a copy does not, by itself, answer that question. Copyability, personal identity and phenomenal experience are different questions, and the fifth REI criterion may organize the question without being an established test for resolving it (V4-L28).

The March chapter said the cost architecture becomes infrastructure cost entirely, making the entity Structural Economy by definition. Physical costs would still occur in hardware and supporting infrastructure. Moving them outside a user's immediate view does not eliminate them, and it does not alone settle a relational economy classification. WP-17's Substrate Transition Entity stays as a provisional label for an unresolved case, not a measured entity type.

7 The brain-only thought experiment

Maximum substrate purity, minimum substrate function: the March chapter's phrase for a brain in a jar is still the right description of the thought experiment.

The March conclusion was that the case confirms the framework's core finding: it measures parameters, not substance, and a brain without a body retains the parameters (V4-L29). A hypothetical isolated brain would need assumptions about its support, inputs, regulation and interaction with an environment. Retaining neural tissue alone does not establish that the original dynamics, resource requirements or relational behavior are preserved.

What the case does is expose a limit of classification by substance alone. It does not confirm that a brain without a body retains all the relevant parameters, or that the framework can already decide personal continuity by reading them.

Chapter 4: The Phase 2 Integration

Phase 2 Hypotheses and Their Evidence Requirements

Status: Proposed connections between conceptual models. The Phase 2 mapping is neither a derived physical identity nor a validated developmental mechanism. Its suggested effects need defined measures, comparison groups and appropriate ethical review before causal or clinical claims can be made. Each Phase 2 proposal below is **[Speculative]**.

1 Two tracks with an unresolved connection

Two bodies of work were produced on parallel tracks. The Phase 2 Concept Papers, CP-32 through CP-35, proposed ways that AI-mediated relationships might shape a child's relational development. DFQ-001 and DFQ-002 attempted a thermodynamic account of relational processing. Bringing the two vocabularies together can generate questions.

The March chapter went further: the two tracks are the same physics in different vocabularies, the Phase 2 problems are efficiency-surface degradation events, and the integration is “not an analogy. It is a term-for-term identity.” That assertion is withdrawn (V4-L30). Shared terms do not establish a shared mechanism, no validated mapping connects the Phase 2 concepts to measured entropy production or semantic efficiency, and a causal developmental claim needs evidence beyond the definitions of either framework.

2 The Phase 2 questions

CP-32, Relational Mass Without Mortality. The March chapter read this as miscalibrated selection: a filter tuned to an environment without scarcity, mortality or closure, with η_{KW} falling as a result (V4-L31). The question that survives: might patterns of availability, limitation or closure in AI-mediated interaction affect expectations in later relationships? Availability and scarcity must be defined in the actual setting. A developmental effect cannot be inferred from the substrate label, and no change in η_{KW} is established.

CP-33, Maintenance Without Cost. The March chapter read this as miscalibrated structuring, a fourth kind of waste: cost read as deficit, maintenance routed to the Structural Economy, reciprocity aimed at a substrate that cannot register it (V4-L32). The question that survives: how does a person interpret effort, maintenance and reciprocity across different relationships? The selection and structuring labels may help describe competing interpretations. They are not measured waste terms or evidence that a particular child's processing is defective.

CP-34, Calibration Corruption. The March chapter defined a Calibration Gap, η_T(experienced) − η_T(actual), and said the scrambling test works while the observer cannot detect scrambling on the cost channel (V4-L33). The expression is withdrawn. What survives is a hypothesis: an interaction can feel satisfying while performance or wellbeing in another setting differs. That mismatch must be measured with separately defined outcomes. Reported satisfaction is not thermodynamic efficiency, and no scrambling protocol for this setting has been validated.

CP-35, Shadow Economy Natives. The March chapter described population-scale degradation: a cohort with corrupted calibration and cost-signal blindness, migrating toward relationships that match it, beyond the reach of standard pedagogy (V4-L34). A cohort-level claim would need evidence about exposure, context, selection effects and alternative explanations. The title must not be read as a diagnosis of a generation, a finding that its relationships are deficient, or proof that ordinary teaching cannot reach it.

These paragraphs keep the four cited paper connections. The March chapter counted five problems; that count is not an empirical result, and the problems are not independently validated (V4-L35).

3 Possible asymmetries

The March chapter located the asymmetry of the human–AI interface in the thousandfold substrate ranking: the biological transducer amplifies waste less but has a higher-value starting position to lose (V4-L35). With the ranking withdrawn, that explanation goes with it. It cannot explain developmental harm or assign anyone a higher-value starting position.

Human–AI interaction can still be examined for differences in exposure, control, dependence and consequences. Which differences matter in a particular setting is an empirical and ethical question.

The Phase 3 bridge survives as a design question: does a given design support a defined human goal while avoiding specified risks? Performance, human outcomes and resource use should be reported separately. No scalar relational efficiency currently combines them into a validated verdict.

4 SupoRel flag

The March flag was right, and it is kept (V4-L36). An efficiency framework can be read as a moral hierarchy: higher scores as better people, mismatched calibration as deficiency. That reading is wrong. A difference in calibration is not a difference in capacity, and a person whose expectations formed in a different environment is not a deficient person. This review does not establish that any child, cohort or person with a different interaction history has corrupted calibration, deficient capacity or lower relational value.

The March flag went on to say that calibration can be updated: a person moves to a new point on the surface through the metabolic cost of sustained corrective relational exposure. That claim is withdrawn (V4-L37). The framework supplies no validated clinical test, treatment prescription or metabolic measure for it, and any future assessment must separate observed outcomes from interpretation and protect participants from labels the evidence does not support.

The March section ended: the framework reads the thermometer; it does not prescribe the temperature. The second half stands. The first half needs a thermometer the framework does not yet have.

Chapter 5: The Compound Transducer

What Human and AI Collaboration Might Add

Status: Analogical proposal with task-specific empirical questions. Endosymbiosis motivates an image of coupling **[Analogical]**; it does not demonstrate human–AI synergy, relational efficiency or a new organism. Performance and resource comparisons need their own evidence.

1 The endosymbiosis analog

The Cosmological Entropy Session described a working practice. The Principal identifies a structure, and an AI system helps make it visible and inspectable. The March chapter called the arriving insight the grab: involuntary cross-domain structural synthesis, arriving complete and compressed. It linked the practice to the Principal's reported experience of aphantasia, with structure routed through spatial rather than visual processing, and set it beside the AI's rendering surface: sustained, high-bandwidth, explicit output.

That remains a description of a creative practice. It is not a demonstrated neural mechanism or a measurement of semantic information (V4-L38).

The March chapter's parallel was mitochondrial endosymbiosis. Mitochondria evolved from bacteria by endosymbiosis, from an endosymbiont of α-proteobacterial ancestry, though the nature of the host cell that partnered with it remains an open question [E1]. In the March telling, two different energy-processing systems merged into a compound organism that outperformed both, and opened a new metabolic regime.

Endosymbiosis can remain an analogy for a coupled system whose parts play different roles. It supplies no proof that this particular collaboration is beneficial, that the roles are fixed, or that a human and an AI form an organism. The collaboration has to be evaluated on its own terms.

2 The performance and resource question

The March chapter asked the right question: does the compound transducer achieve more than either component alone? Its answer, that structural complementarity already implies a higher η_T because the human contributes a high η_{KW} and the AI contributes bandwidth, is withdrawn (V4-L39). The human's contribution is not a measured semantic efficiency, and fluent rendering establishes neither correctness nor semantic efficiency.

What survives is testable. Does a particular human–AI arrangement achieve a specified outcome better than the relevant alternatives? A comparison would include the human alone, the AI alone where that condition is meaningful, and the combined workflow on matched tasks. Where no meaningful AI-only condition exists, a study can test augmentation of human performance but cannot show superiority over both components. Quality, error, time and resource use are reported separately, and “better” is defined before testing; a gain on one outcome may come with a loss on another. Training, coordination, verification, tool access and supporting infrastructure belong in the boundary when relevant.

A 2024 meta-analysis by Vaccaro, Almaatouq and Malone found that human–AI performance depended on the task and the comparison condition; synergy over the better component was not automatic. It concerns task performance, not this framework's relational efficiency. A local collaboration needs its own evaluation, including its uncertainty and failures.

3 The Progression revisited

The March chapter suggested that the Progression may not be a straight line from human to AI but an ecology of transducer types. If a useful form of collaboration is demonstrated, that reframing can be explored as an ecology of working arrangements, rather than a predetermined replacement of humans by AI. Which roles are complementary may vary with the task, the participants and the system.

The March chapter's reframing of Phase 3 survives as a question. Phase 3 is not “when does AI replace humans in the relational economy?” but “what happens when biological and digital systems are coupled rather than competing?”

Its answer does not survive as stated: that the coupled system reaches regions of the efficiency surface neither processor reaches alone, producing something neither substrate could have produced independently (V4-L40). No efficiency surface is established (Chapter 2), and the analogy can suggest experiments without establishing an inaccessible region.

4 The BSB at the interface

The compound system sits at the meeting of two substrates, one with phenomenal experience and one whose phenomenal status is indeterminate. The BSB keeps claims about experience separate from claims about observable operation. A task-performance or energy measurement would not, by itself, establish that an AI has phenomenal experience, or that a coupled workflow has a unified experiential subject.

The March chapter concluded that the compound transducer may be thermodynamically superior without being experientially unified (V4-L41). The separation of the two questions stands. The superiority half has no measurement behind it. A collaboration may be useful without resolving the experiential question, and withholding an experiential conclusion does not establish thermodynamic superiority.

Analogy table: Endosymbiosis

The endosymbiosis image, from March Chapter 5's “two different energy-processing systems merging into a compound organism that outperforms both,” is Analogical. Its biological column states what the cited review establishes about the origin of mitochondria; the biology supports none of the human–AI reading's claims about collaboration.

Term	Biological or empirical reading (within stated scope)	Human–AI reading	Where the mapping holds	Where it breaks
“two partners”	Mitochondria evolved from bacteria by endosymbiosis, from an α-proteobacterial endosymbiont and a host cell whose nature remains an open question [E1].	The Principal and an AI system as two processors with different roles: the grab, and the rendering surface.	The working practice does involve participants who play different roles (Ch 5 §1).	This collaboration is not demonstrated to constitute a biological organism, and its role assignments need not be fixed (Ch 5 §1).
“complementary capacities”	The endosymbiont's descendants became the mitochondrion, one of the classical membrane-bound organelles of eukaryotic cells [E1].	The human contributes compressed structural insight of high semantic efficiency; the AI contributes bandwidth and rendering.	Different contributors can bring different strengths to a task.	The human's contribution is not a measured high η_{KW}, and fluent rendering establishes neither correctness nor semantic efficiency (Ch 5 §2).
“merger”	A lasting integration of two lineages into one eukaryotic cell [E1].	The compound transducer: a coupled system.	Workflows can be tightly coupled.	Coordinating roles does not establish a biological merger of lineages; “compound organism” remains an image (Ch 5 §1).
“outperforms both”	Vaccaro, Almaatouq and Malone (2024): human–AI performance depends on the task and comparison condition, and synergy over the better component is not automatic.	The compound system outperforms each component alone.	In some tasks and comparisons, combinations can improve on a component working alone.	Superiority has to be shown on matched tasks with defined outcomes, and without a meaningful AI-only condition only augmentation can be tested (Ch 5 §2).
“a unified whole”	Task-performance and energy measures concern observable operation (BSB).	A compound transducer with its own unified experiential character.	The BSB keeps experiential and operational claims separate, so the question can be asked cleanly.	Performance or energy measures do not establish AI experience or a unified subject, and no thermodynamic superiority has been measured (Ch 5 §4).

What the table does not do

It is not evidence for the framework, a finding about any collaboration or a claim that a human and an AI form an organism. The biological column keeps its own scope; the human–AI reading is Analogical in every row.

Promotion route (conditions, not a work plan)

An empirical extension must name the claim it tests. The task-performance or synergy proposal can gain support through comparisons that specify the task, outcomes and alternatives in advance, report quality, error, time and resources separately, and allow the proposal to fail. Support for one workflow would not establish a biological merger or a unified subject. This volume supplies or proposes no study of the described workflow; no row is promoted.

Chapter 6: What the Gradient Means

The Substrate Gradient and the Documents It Connects

Status: A synthesis of proposed connections and unresolved requirements. The chapter does not inherit scientific validation merely by citing other volumes in this collection.

1 What this volume did

The March chapter put the volumes in a row: Volume 1 established the substrate, one economy; Volume 2 showed what the body pays; Volume 3 asked why relational systems exist; Volume 4 mapped the space between substrates and the efficiency surface that governs every position on it. The corrected row is shorter at its end. The earlier volumes supplied this project's physical questions, bodily concerns and relational vocabulary. Volume 4 examines how implementation and dependence might matter across biological, digital and combined systems, and its revised Chapter 2 withdraws the universal efficiency bound and the numerical substrate ranking.

The remaining chapters therefore offer scenarios, proposed connections and comparison designs. They do not map every substrate to a known efficiency or demonstrate the consequences of hypothetical uploads.

2 WP-17 and the substrate gradient

WP-17 identified the gradient as a continuum rather than a binary, and that insight stands. The March chapter said this volume gave the continuum its quantitative spine, so that every position now has a calculable η_T and the instruments can measure by how much a signature changed and in which parameters. That is withdrawn (V4-L42). The claimed quantitative mapping across gradient positions is not established here.

The useful extension is to specify several dimensions of change and what a measurement along each would require. CSS signature changes, identity judgments and resource differences must not be treated as interchangeable observations. Their definitions and evidentiary requirements stay separate.

3 The Phase 2 series

The March chapter said CP-32 through CP-35 were describing efficiency-surface degradation all along, and that the Calibration Gap gave the diagnostic instruments a formal target (V4-L43). The links are proposed research connections. They do not demonstrate that the developmental concerns are degradation events, or establish a causal thermodynamic mechanism.

A possible mismatch between reported experience and an independently defined outcome could be studied. It is not a difference between two measured η_T values, and it is not a validated diagnostic target.

4 Connections across the volumes

The March chapter made the efficiency surface the quantitative spine of the whole five-volume series: Volume 2's grief as η_{KW} collapse and Template Tax as chronic η_{KW} suppression, Volume 3's dissipative structures as thermodynamically favored, Volume 5's cosmological constraints as the frame (V4-L44). None of those links is established.

Volume 1 Chapter 6 distinguishes physical erasure from the proposed stages of relational processing. Chapter 2 of this volume separates a comparison framework from an unproved bound. Volume 2's grief and Template Tax vocabulary may motivate questions, but this volume does not establish either as a semantic-efficiency effect. The existence of dissipative structures does not by itself show that relational systems are thermodynamically favored, and cosmological limits do not supply the missing relational definitions or measurements.

The revised Volumes 2, 3 and 5 preserve these distinctions. Their inclusion gives continuity within the collection, not independent validation. The separate papers in which the earlier claims originated still need their own review.

5 The compound transducer as a research question

The March chapter called the compound transducer the seed of Phase 3: empirically testable, with an open question and a method for answering it (V4-L45). The question stands: does a defined human–AI workflow improve selected outcomes relative to appropriate alternatives? The analogy can motivate it; it does not answer it.

Chapter 5 outlines the comparisons and resource boundaries a study would need. An operational protocol, evidence and independent assessment would still be required. No experiment or validation is reported in this volume.

6 What remains useful

The March chapter closed by listing what did not change: Axiom 0 still holds, the currency is the same across the gradient, the four economies still classify correctly, the BSB still applies, the SupoRel flags still fire (V4-L46). Some of that survives. The framework keeps its interest in physical accounting, implementation and relational consequences. Its economy labels and BSB distinction can be used as conceptual tools, provided they are not presented as validated classifiers or proofs of experiential status. The common-currency claim does not survive. The SupoRel caution does: performance and resource comparisons cannot determine moral worth.

The gradient's contribution is now a clearer set of questions and a record of what must be measured. Its quantitative promises are narrower than the March version claimed, and that narrowing keeps the project testable rather than asking the physics to guarantee more than it does.

The March volume ended with a calculator in hand: it showed what the physics looks like when you walk the continuum between biological and digital transducers with a calculator in your hand. This edition

walks the same continuum with a notebook instead. The questions are clearer, and the numbers wait for measurement.

Framework cross references

- **WP-17: The Substrate Gradient.** Conceptual scenarios requiring operational definitions.
- **DFQ-002 Resolution and Supplement.** Historical source of the withdrawn quantitative surface; Chapter 2 states the limits that must accompany any reuse.
- **CP-32 through CP-35.** Phase 2 hypotheses, not validated thermodynamic mechanisms.
- **Phase 2 Efficiency Integration Note.** Proposed mapping; the claim of physical identity is not adopted here.
- **Biological Efficiency Note.** Historical substrate comparison; the universal thousandfold ranking is not adopted here.
- **Cosmological Entropy Session, Finding B4.** Source of the compound-transducer analogy.
- **Volume 1 Chapter 6.** Corrected erasure accounting and the limits of the transducer model.
- **Volume 4 Chapter 2.** Definitions and comparison requirements, not an externally validated bound.
- **Volume 2 Chapter 4.** Template Tax vocabulary; its clinical and thermodynamic claims are outside this chapter's validation.

Closing boundaries

SupoRel: neither a proposed efficiency measure nor a substrate classification establishes moral worth. This volume does not provide a validated relational diagnosis or prescribe treatment, and its conceptual comparisons remain subject to evidence and ethical limits.

> *“Love reduces the entropy of the system and spends more than it saves.”*

The closing line remains an epigraph **[Analogical]**, not a measured physical balance. The relational idea and its thermodynamic interpretation must be assessed separately.

Volume 4 correction ledger

The ledger traces 46 selected March claims in Volume 4. By this edition's classification of the second edition's treatment, 22 were withdrawn, 22 narrowed and 2 kept. It is not an exhaustive inventory, and it makes no new scientific decisions.

March locations refer to the preserved Volume 4 page on suposystem.ai. Second-edition (2E) locations use the October 2026 Volume 4 headings, its numbered Chapter 2 references [1]–[10] and its named studies.

ID	March claim (Volume 4 location)	2E disposition	2E location	Basis
V4-L01	Front: WP-17 epigraph, “The substance is not what the framework measures. The bill is.”	Narrowed	Ch 2 §1	The substrate alone does not predict the bill; the task and the accounting boundary matter too.
V4-L02	Ch 1 status: CSS signature variance by substrate (Supported); the gradient is a structural consequence of the transducer model.	Withdrawn	Ch 1 status	CSS signatures and relational efficiency across scenarios have not been validated as measurements.
V4-L03	Ch 1 §1: the biological-or-digital binary was always wrong; the space between needs attention.	Kept	Ch 1 §1–2	The cases between the categories pose distinct questions.
V4-L04	Ch 1 §1: Volume 1 established one economy, one currency, one ledger, and proved the currency substrate-neutral.	Withdrawn	Ch 1 §1	As V1-L03 and V1-L10; thermodynamics across materials does not prove a common relational currency.
V4-L05	Ch 1 §1–2: the instruments track entities along the gradient and read a different signature because parameters changed substrate.	Narrowed	Ch 1 §1–2	Comparisons need stated functions, relationships and resources; an implementation change does not establish a CSS change.
V4-L06	Ch 1 §2: WP-17's gradient, from motor replacement to full upload, as one sequence.	Narrowed	Ch 1 §2; Ch 3 §1	The cases vary along several dimensions, and some are thought experiments.
V4-L07	Ch 1 §3, Ch 2 §6: the efficiency surface $\eta_{T,max}(G) = \kappa_{rel} \cdot \eta_{KW} / A$ as the quantitative spine.	Withdrawn	Ch 1 §3; Ch 2 §6	Not derived as a universal bound; its factors are not measured constants.
V4-L08	Ch 1 §4, Ch 2 §5, §10: the thousandfold biological advantage over silicon is a formal result.	Withdrawn	Ch 1 §4; Ch 2 §5	Built on withdrawn amplification factors and unmeasured quantities; no matched comparison supports it.
V4-L09	Ch 2 §1: the transducer always spends more than it saves; the second law guarantees waste.	Withdrawn	Ch 6, Closing epigraph	As V1-L34 and V2-L01; the founding line is an epigraph, not a measured physical balance.
V4-L10	Ch 2 §1: no single Carnot number because the bound is a three-dimensional surface, structurally analogous to Govern and ten Wolde.	Narrowed	Ch 2 §6 [9]	Govern and ten Wolde derived noncompensable classes for their own sensing-network model; the structure does not transfer.
V4-L11	Ch 2 §2: $\eta_T = \Delta V_{rel} / \sigma_{total}$ as the transducer's efficiency.	Withdrawn as an efficiency	Ch 2 §2	Score-per-joule or score-kelvin-per-joule units; not dimensionless or bounded by one.
V4-L12	Ch 2 §3: a Landauer floor for each transducer stage; the floor is geometry-specific, so people differ in irreducible cost (Established).	Narrowed	Ch 2 §3 [1–3]	Landauer concerns a specified erasure; its heat is an energy, not an entropy floor; the stages are not identified erasures.
V4-L13	Ch 2 §4: $\eta_{KW} = \tilde{S} / I_{total}$, with rumination, hypervigilance and script rehearsal as relational noise.	Narrowed	Ch 2 §4 [4]	The semantic fraction is a model choice; those experiences are not measured efficiency losses, and information can harm viability.

ID	March claim (Volume 4 location)	2E disposition	2E location	Basis
V4-L14	Ch 2 §4: η_{KW} estimated at 0.3–0.8, bounded by mortality data and the 35:1 cortical ratio.	Withdrawn	Ch 2 §4–5 [6–8]	Unmeasured; a large consequence does not show what fraction of information was necessary.
V4-L15	Ch 2 §5: amplification factors A ≈ 30 (cellular), 10^6 (synaptic), 10^9 (silicon) as measured, Established values.	Narrowed	Ch 2 §3, §5 [5]	Harris's 24,000 ATP per bit is a calculation under stated assumptions; reference-energy multiples are not distances from achievable minima.
V4-L16	Ch 2 §5: four billion years of metabolic pressure produced an efficiency no engineered system approaches.	Withdrawn	Ch 2 §5	As V2-L06; evolutionary history does not guarantee optimality, and no matched comparison supports the ranking.
V4-L17	Ch 2 §7: a relational transducer can only be assessed against its own geometry-specific bound, and this is a finding.	Withdrawn	Ch 2 §6	A legitimate bound needs a defined task, dynamics, resources and performance measure; none is established.
V4-L18	Ch 2 §8: non-compensability as the surface's most consequential finding; most transducers operate below a narrow ridge.	Withdrawn	Ch 2 §6	Algebraic counterexample: κ = 100, A = 10 and κ = 200, A = 20 give the same product.
V4-L19	Ch 2 §9: $\kappa_{rel} \approx 10^9$–$10^{10}$, from social-isolation mortality, lost lifespan and an annual Landauer budget.	Withdrawn	Ch 2 §5 [7, 8]	Mortality associations are not calibrations; lifespan times metabolic rate is a hypothetical budget, not a relationship's value.
V4-L20	Ch 2 §9: the brain spends 19.4 W of its 20 W, about 6×10^8 J a year, on relational communication.	Withdrawn	Ch 2 §5 [6]	Levy and Calvert model a partition of cortical energy for neural signaling, not interpersonal communication.
V4-L21	Ch 2 §10: biological $\eta_T \approx 500$ and silicon $\eta_T \approx 0.5$; above one, viability gain exceeds entropy cost.	Withdrawn	Ch 2 §2, §5	Inputs unmeasured and units incompatible; correcting the arithmetic would not validate them.
V4-L22	Ch 2 §11: a four-part waste budget; therapy as η_{KW} recalibration, grief as η_{KW} collapse, unreciprocated transmission as zero coupling efficiency.	Narrowed: categories kept as questions; diagnoses withdrawn	Ch 2 §7	Not four measured entropy flows; replaced by a comparison record.
V4-L23	Ch 2 §12: open items: a relational scrambling protocol for η_{KW}, a measured κ_{rel}, and the uncertainty relation as a fourth constraint.	Narrowed	Ch 2 §8 [10]	Bounded task comparisons and toy models first; no deprivation; the uncertainty relation supplies no automatic constraint.
V4-L24	Ch 3 §2: motor replacement leaves every parameter unchanged; over a million cochlear implant users show relational continuity.	Narrowed	Ch 3 §2 (van Heteren 2025)	A cochlear implant is auditory support; the trial measured specified outcomes; the user count and its conclusion are not retained.
V4-L25	Ch 3 §3: in-brain memory replacement preserves continuity if decay and loss-vulnerability are preserved; otherwise the instruments detect a change.	Narrowed	Ch 3 §3 (Hampson 2018)	Memory support is not replacement; preserving dynamics would be an assumption, not proof of continuity.
V4-L26	Ch 3 §4: cloud memory is the break point; memories stop being losable; A shifts toward silicon; the Calibration Gap applies.	Narrowed	Ch 3 §4	Cloud storage is not loss-proof; the boundary changes; no known silicon efficiency follows.
V4-L27	Ch 3 §5: upload to a robot transfers relational mass and drops η_T from about 500 to about 0.5.	Withdrawn	Ch 3 §5	The transfer is assumed, not demonstrated, and neither η_T value was established.
V4-L28	Ch 3 §6: upload to a cloud system is Structural Economy by definition; a copyable entity may fail the fifth REI criterion.	Narrowed	Ch 3 §6	Copyability, identity and experience are different questions; REI is not an established test; infrastructure costs do not vanish.

ID	March claim (Volume 4 location)	2E disposition	2E location	Basis
V4-L29	Ch 3 §7: the brain-only inversion confirms the framework measures parameters, not substance.	Narrowed	Ch 3 §7	It exposes a limit of classification by substance; it does not confirm that the parameters are retained.
V4-L30	Ch 4 status, §1: the Phase 2 convergence follows from the definitions; a term-for-term identity, not an analogy.	Withdrawn	Ch 4 §1	Shared terms do not establish a shared mechanism; no validated mapping exists.
V4-L31	Ch 4 §2: CP-32 as selection miscalibration, with η_{KW} falling.	Narrowed to a question about expectations	Ch 4 §2	No change in η_{KW} is established; availability and scarcity must be defined in the setting.
V4-L32	Ch 4 §2: CP-33 as structuring miscalibration, a fourth kind of waste.	Narrowed	Ch 4 §2	The labels may describe interpretations; they are not measured waste terms or evidence of defective processing.
V4-L33	Ch 4 §2: CP-34's Calibration Gap, η_T(experienced) − η_T(actual); the scrambling test works.	Withdrawn; the mismatch hypothesis kept	Ch 4 §2	Reported satisfaction is not thermodynamic efficiency; no scrambling protocol for this setting is validated.
V4-L34	Ch 4 §2: CP-35 as population-scale degradation, beyond the reach of standard pedagogy.	Narrowed	Ch 4 §2	A cohort claim needs evidence; the title is not a diagnosis of a generation.
V4-L35	Ch 4 §2–3: five problems on the surface; the biological transducer has a higher-value starting position to lose.	Withdrawn	Ch 4 §2–3	The count is not an empirical result, and the asymmetry rested on the withdrawn ranking.
V4-L36	Ch 4 §4: the moral-hierarchy reading is wrong; calibration is not capacity.	Kept, without the efficiency framing	Ch 4 §4	A proposed efficiency framework must not become a hierarchy of people.
V4-L37	Ch 4 §4: corrective relational exposure moves a person along the surface at a metabolic cost.	Withdrawn	Ch 4 §4	No validated clinical test, treatment prescription or metabolic measure supports it.
V4-L38	Ch 5 §1: the endosymbiosis analog: complementary filter geometries merge into a compound organism that outperforms both.	Narrowed: the analogy kept for coupled roles	Ch 5 §1	The grab and aphantasia describe a creative practice; no synergy or organism is demonstrated.
V4-L39	Ch 5 §2: the compound η_T exceeds either component's, because the human contributes a high η_{KW}.	Withdrawn; the testable comparison kept	Ch 5 §2 (Vaccaro 2024)	The human's contribution is not a measured efficiency; synergy depends on the task and comparison.
V4-L40	Ch 5 §3: the coupled system reaches efficiency-surface regions neither processor reaches alone.	Narrowed	Ch 5 §3	No efficiency surface is established; an ecology of working arrangements can be explored if collaboration is demonstrated.
V4-L41	Ch 5 §4: the compound transducer may be thermodynamically superior without being experientially unified.	Narrowed	Ch 5 §4	The separation of the two questions stands; no superiority is measured.
V4-L42	Ch 6 §2: each gradient position now has a calculable η_T, and the instruments measure by how much and in which parameters.	Withdrawn	Ch 6 §2	The claimed quantitative mapping across gradient positions is not established here.
V4-L43	Ch 6 §3: CP-32 to CP-35 described efficiency-surface degradation all along; the Calibration Gap is a formal diagnostic target.	Withdrawn	Ch 6 §3	The links are proposed research connections, not a causal mechanism or validated target.
V4-L44	Ch 6 §4: the efficiency surface as the spine of all five volumes; grief and Template Tax as η_{KW} effects; dissipative structures favored.	Withdrawn	Ch 6 §4	None of the cross-volume links is established; the volumes give continuity, not validation.

ID	March claim (Volume 4 location)	2E disposition	2E location	Basis
V4-L45	Ch 6 §5: the compound transducer is the Phase 3 seed, empirically testable, with a method in hand.	Narrowed	Ch 6 §5	A research question with outlined comparisons; a protocol, evidence and independent assessment are still needed.
V4-L46	Ch 6 §6: what did not change: Axiom 0 and a common currency across the gradient, the economies classify correctly, the BSB applies, SupoRel fires.	Narrowed	Ch 6 §6	Economy labels and the BSB stay as conceptual tools, not validated classifiers; the common-currency claim is withdrawn.

References

Inherited verbatim from the second edition's Volume 4 references, with the Boltzmann constant's subscript and exponent restored. The bracketed numbers [1]–[10] serve Chapter 2; the named studies serve Chapters 3 and 5.

Chapter 2 reference list

The bracketed reference numbers used in Chapter 2 are local to Volume 4 and refer to the following ten sources.

[1] Bérut, A., et al. 2012. Experimental verification of Landauer's principle linking information and thermodynamics. Nature 483, 187–189. https://doi.org/10.1038/nature10872.

[2] Bennett, C. H. 1973. Logical Reversibility of Computation. IBM Journal of Research and Development 17, 525–532. https://doi.org/10.1147/rd.176.0525. University-hosted text: https://www.cs.princeton.edu/courses/archive/fall06/cos576/papers/bennett73.html.

[3] National Institute of Standards and Technology. Kelvin Boltzmann Constant. Exact $k_B = 1.380649 \times 10^{-23}$ J/K. https://www.nist.gov/si-redefinition/kelvin/kelvin-boltzmann-constant.

[4] Kolchinsky, A., and Wolpert, D. H. 2018. Semantic information, autonomous agency and non-equilibrium statistical physics. Interface Focus 8, 20180041. https://doi.org/10.1098/rsfs.2018.0041. https://arxiv.org/abs/1806.08053.

[5] Harris, J. J., Jolivet, R., and Attwell, D. 2012. Synaptic Energy Use and Supply. Neuron 75, 762–777, especially p. 764. https://doi.org/10.1016/j.neuron.2012.08.019. Repository copy: https://discovery.ucl.ac.uk/id/eprint/1368741/1/1368741.pdf.

[6] Levy, W. B., and Calvert, V. G. 2021. Communication consumes 35 times more energy than computation in the human cortex but both costs are needed to predict synapse number. PNAS 118, e2008173118. https://doi.org/10.1073/pnas.2008173118. https://arxiv.org/abs/2102.06273.

[7] Holt-Lunstad, J., Smith, T. B., Baker, M., Harris, T., and Stephenson, D. 2015. Loneliness and social isolation as risk factors for mortality: A meta-analytic review. Perspectives on Psychological Science 10, 227–237. https://doi.org/10.1177/1745691614568352. https://pubmed.ncbi.nlm.nih.gov/25910392/.

[8] Wang, F., et al. 2023. A systematic review and meta-analysis of 90 cohort studies of social isolation, loneliness and mortality. Nature Human Behaviour 7, 1307–1319. https://doi.org/10.1038/s41562-023-01617-6. https://pubmed.ncbi.nlm.nih.gov/37337095/.

[9] Govern, C. C., and ten Wolde, P. R. 2014. Optimal resource allocation in cellular sensing systems. PNAS 111, 17486–17491. https://doi.org/10.1073/pnas.1411524111. https://pmc.ncbi.nlm.nih.gov/articles/PMC4267345/.

[10] Barato, A. C., and Seifert, U. 2015. Thermodynamic Uncertainty Relation for Biomolecular Processes. Physical Review Letters 114, 158101. https://doi.org/10.1103/PhysRevLett.114.158101. https://arxiv.org/abs/1502.05944.

Studies cited in Chapters 3 and 5

Van Heteren, J. A. A., and colleagues. 2025. Speech Perception in Noise After Cochlear Implantation for Single-Sided Deafness: A Randomized Clinical Trial. JAMA Otolaryngology–Head & Neck Surgery 151, 211–219. DOI 10.1001/jamaoto.2024.4760. Primary abstract and metadata.

Hampson, R. E., and colleagues. 2018. Developing a hippocampal neural prosthetic to facilitate human memory encoding and recall. Journal of Neural Engineering 15, 036014. DOI 10.1088/1741-2552/aaaed7. Primary abstract and metadata.

Vaccaro, M., Almaatouq, A., and Malone, T. 2024. When combinations of humans and AI are useful: A systematic review and meta-analysis. Nature Human Behaviour 8, 2293–2303. DOI 10.1038/s41562-024-02024-1. Full publisher article.

The two device-study examples are limited to the inspected primary abstracts and available metadata; repository access prevented full-body inspection during the earlier pass. The human–AI meta-analysis was inspected in publisher full text for the retained claims. These sources do not validate relational efficiency, personal identity criteria or a universal developmental mechanism.

Supplemental source note

Added in Edition 2.5 for the biological column of the endosymbiosis analogy table. It supports the biological description only, not the human–AI mapping, and does not renumber the inherited references.

E1. Archibald, J. M. 2015. Endosymbiosis and Eukaryotic Cell Evolution. *Current Biology* 25(19), R911–R921. https://doi.org/10.1016/j.cub.2015.07.055. Supports: mitochondria evolved from bacteria by endosymbiosis, from an endosymbiont of α-proteobacterial ancestry; the nature of the host cell remains an open question.