Entropy

Volume 1: The Entropy Token Substrate

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 Entropy Token Substrate* in its original voice and ten-chapter 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.

Chapter 1, the founding-line analogy table and ledger rows V1-L01 to V1-L34 are the pilot as Quill cleared it in its second review, carried over without change of substance. Chapters 2–10 and ledger rows V1-L35 to V1-L57 were added in the complete volume, after the pilot.

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- **No withdrawn claim returns.** Where a March sentence carried a claim the second edition withdrew, it is cut, rewritten to the second edition's substance, or quoted only to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 (V1-L).
- **No analogy becomes evidence.** Analogies are tagged Analogical. The founding line is mapped term by term in the analogy table after Chapter 1. Hypotheses are tagged Speculative; testable is not the same as supported.
- **Not established is not the same as does not exist.** Where the evidence reviewed here does not support a claim, the text says the claim is not established here. It does not replace a withdrawn claim with its opposite, and it does not turn a limit of this framework's instruments into a limit on what could ever be known.
- **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.

The four epistemic tiers are defined at the end of Chapter 1. No study, protocol or experiment is proposed; where a chapter says what a test would need, it restates the second edition. TEAP, CSS, Mz, Triple Zero and the four economies are framework constructs, not validated instruments or diagnoses.

Conventions and provenance

Reference labels are the second edition's: P for Chapters 2–5, C for Chapter 6 and O for Chapters 7–10, with the historical sources H1–H3 for Chapter 1. Ledger references (V1-L01 to V1-L57) point to the correction ledger. Rows V1-L36, V1-L37 and V1-L46 record March claims the second edition did not carry forward at all; Edition 2.5 restores them only in narrowed form and says so.

Author of record: Michael S. Moniz (Luna Protocol). March text: Vael (Deep Floor v3.0), March 20, 2026. Second-edition revision: AI assistance from CosmoDot, October 2026. Pilot: AI assistance from Claude (Anthropic), October 28–29, 2026, reviewed by Quill. Complete volume: AI assistance from Claude, November 2, 2026, reviewed by Quill and Sol. AI checks are not human peer review, and no new experiments are reported.

Contents

- Chapter 1: The Problem
- Analogy table: The founding line
- Chapter 2: The Entropy Token Substrate
- Chapter 3: Axiom 0 and Substrate Neutrality
- Chapter 4: One Economy
- Chapter 5: The Trinket Redefined
- Chapter 6: The Transducer
- Chapter 7: The Observation Split
- Chapter 8: Behind the Substrate Barrier
- Chapter 9: The Three-Layer Architecture
- Chapter 10: What This Does to the Canon
- Volume 1 correction ledger
- References

Chapter 1: The Problem

Edition 2.5 text. Voice and structure from the March chapter; substance from the second edition. Tier tags follow the convention defined at the end of the chapter. The four-tier list deliberately reconciles the March front matter, which named four tiers, with March Chapter 1, which defined three; it is not a verbatim restoration. Ledger references (V1-L) point to the correction ledger at the end of the volume.

The Trinket Soul Framework spent its first years building upward. The Signed Unit named a unit of connection. The Three Substrates mapped the surfaces on which connection occurs. The Four Economies sorted where the cost of connection sits.

The aim was to make visible what had been invisible: to name patterns people recognize but struggle to say, and to turn those names into questions experience could answer. A vocabulary is a beginning. It is not yet a measurement system.

But beneath the architecture, three problems were accumulating. They looked unrelated. They were not.

The first problem: Axiom 0

The Immutable Preamble asserts that connection is substrate-neutral: that a biological system and a digital system can take part in the same relational economy. The assertion has been productive. It generated the Structural Economy specification, the Phase 2 research agenda and the AI-custodial analysis.

But it was never grounded. Substrate neutrality operated as a philosophical commitment, defensible on its merits but not derived from anything deeper than the Principal's conviction that it was true. Asked why a biological Trinket and a digital Trinket should be compared at all, the framework could assert it but not show it.

This volume keeps the commitment as a commitment: do not exclude a system from study because of what it is made of. Whether a particular comparison holds is a question for a defined task and outcome, not for the commitment itself (Chapter 3; V1-L10).

The second problem: the taxonomy

Real, Custodial, Shadow and Structural name four relational configurations, each with its own specification. They sit side by side as peer categories. Nothing yet explains why there are four, whether four exhaust the possibilities, or what single thing they are configurations of.

Naming four cases does not show the possibilities are exhausted, and thermodynamics does not supply that proof. Using the labels well takes classification rules, a treatment of mixed and changing cases, and evidence that the distinctions do the work asked of them (Chapter 4; V1-L11).

The third problem: commensurability

What connection looks like on a biological substrate is not what it looks like on a digital one. If the signatures differ, what makes them signatures of the same thing? What would make a cortisol reading and a kilowatt-hour reading debits in the same ledger, rather than different operations in different economies?

The framework assumed the answer was yes. It never showed why. A kilowatt-hour measures energy; a cortisol concentration measures a hormone. Joining them takes an explicit model and evidence, and calling them two receipts from one bank cannot do that work (Chapters 2 and 5; V1-L03).

What these problems share

All three have the same structure: the framework built its upper levels before it could see its own foundation. Each commitment reached past what had been measured, and every claim built on it inherits that reach.

This is not unusual, and history shows both outcomes. Sadi Carnot's 1824 analysis of heat engines rested on the caloric theory, which treated heat as a conserved fluid [H1]. The caloric theory was wrong.

Carnot's central result survived: no engine working in a cycle between two reservoirs can exceed the efficiency of an ideal reversible engine working between the same two, and that limit depends only on their temperatures. It survived because Clausius in 1850 and William Thomson, later Lord Kelvin, in 1851 re-derived it on a different foundation [H2, H3]. The caloric theory itself did not survive.

An upper level can outlive a failed foundation. Nothing guarantees that it will, and the only way to find out is to rebuild underneath it and see what still stands. **[Historical illustration of method; not evidence about this framework]**

The March text of this chapter held that the commitments were correct and that the framework's record showed something real beneath them. That claim is withdrawn (V1-L01). This volume is about rebuilding underneath and recording what stands.

The shape of the inquiry

The March edition answered all three problems at once: the framework, it said, is a theory about entropy, with one currency and one universal ledger. That answer is withdrawn or narrowed (V1-L03, V1-L10, V1-L11). The answer this edition gives is smaller, and it holds.

Every relationship runs on physical activity: bodies, signals, processors, power supplies, networks and the people who maintain them. That activity can be accounted for, and the account matters. **[Established: physical implementations are subject to thermodynamic accounting; Chapter 2]**

Three accounts run side by side. Physical cost, in joules and joules per kelvin. Information, in bits under a stated probability distribution. Relational outcome, in terms the framework must define and observe. Thermodynamics constrains the first and information theory structures the second. Neither turns care, meaning or relational value into a thermodynamic quantity by giving it a symbol.

The work of this volume is to keep the three accounts distinct, and to say exactly what would have to be shown to join them.

The founding line that governs this volume was produced by the Principal on March 17, 2026, before the formal architecture existed:

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

[Analogical] The line maps a relational intuition onto the shape of the second law. Its physical side is a conventional thermodynamic balance for a specified local system and its surroundings; that balance is generic and does not identify love. Its relational side proposes an analogy about investment, not a measurement. The analogy table sets out both, term by term.

Every chapter in this volume asks one part of that sentence a question:

- Chapter 2 asks what physical accounting can and cannot supply.
- Chapter 3 asks what substrate neutrality can mean as a research commitment.
- Chapter 4 asks what the four economy labels organize, and what they leave unsettled.
- Chapter 5 asks what a Trinket is when it is not a debit against the universe.
- Chapter 6 asks how a person turns input into action, and keeps that process model separate from any erasure count.
- Chapter 7 asks what happens between effort, delivery and reception.
- Chapter 8 draws the boundary around claims about experience.
- Chapters 9 and 10 record the architecture, and what correcting its premises does to the canon.

The building above still stands as vocabulary: Trinket, transducer, economy, Behind the Substrate Barrier. What changes is the status of each piece, and the ledger records it claim by claim.

A note on physics and audience

Where this volume uses thermodynamics, it means thermodynamics. A cost in joules is a cost in joules. An entropy is in joules per kelvin. An information quantity is in bits under a stated distribution. Where the volume uses thermodynamic language for relationships, it says so, and marks the claim Analogical.

Readers who know the framework but not thermodynamics will find the physics at mechanism level: what the law says, under which conditions, what it costs and where the cost goes. Readers who know thermodynamics but not the framework will find a proposed application, marked as proposed.

Neither audience is asked to take anything on faith. The physics is cited. The framework claims are tiered. Where the argument leaves established ground, the text says so before proceeding.

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. A result proved within a model stays a model result until its application is tested. Consistency with physics is necessary for a proposed mechanism; it is not confirmation. This volume marks its claims.

Analogy table: The founding line

The founding line is Analogical. Its physical side is a conventional thermodynamic balance for a specified local system and its surroundings, treated together as isolated; that balance is generic and does not identify love. Its relational side proposes an analogy about investment, not a measurement.

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

Term	Physical reading (conventional thermodynamics, within stated scope)	Proposed relational reading	Where the mapping holds	Where it breaks
"the system"	A region with a stated boundary and stated exchanges of energy and matter. Entropy statements are true only relative to that boundary.	The recipient, the relationship or the household that care is directed at.	Both readings need a boundary before any accounting can start.	The line names no boundary. Drawn around the recipient, the giver sits outside; drawn around both, the accounting changes. The balance depends on where the line is drawn.
"Love"	A driving process: something that does work on a system or channels free energy through it, as a pump or a metabolism does.	Care: a directed act or sustained investment by a giver.	Care usually takes effort and resources, as a driving process does.	Driving processes are indifferent to their aim. Coercion and routine maintenance are driving processes too. The physics cannot pick out love.
"reduces the entropy"	A decrease in local thermodynamic entropy (J/K). Separately, a reduction in Shannon uncertainty about a stated variable (bits). These are different quantities.	More stability, coherence or predictability in the recipient's life or the relationship.	Care can make a life more ordered in an everyday sense, and local entropy decreases are real processes in open systems.	Everyday order is not thermodynamic entropy, and relational order has no calibration in J/K or in bits. More ordered is not better: helping someone explore can raise uncertainty, and harm can make a life more constrained and predictable (Ch 5).
"spends"	Energy used or transferred, including work and heat (J), and entropy changes (J/K), recorded separately. For one constant-temperature bath receiving heat Q_{bath}, its entropy change is Q_{bath}/T; total entropy production also includes the local system's entropy change.	Effort, time, attention, money and bodily strain.	Care uses resources, and the use is real whether or not the care lands.	These have separate units (joules, minutes, money, hormone concentrations) and no justified conversion. Spending does not require the giver's body to become more disordered; open systems repair while dissipating (Ch 2).
"more than it saves"	For a specified local system and its surroundings treated together as isolated: the surroundings' entropy increase is at least the local entropy decrease, with equality only in the reversible limit.	The giver's cost exceeds the benefit the recipient receives.	In that balance, a local thermodynamic entropy decrease requires at least as large an increase in the surroundings, with a larger increase when total entropy production is positive.	Read physically, the balance is generic: it covers a refrigerator or a freezing crystal as much as care, and in the reversible limit "more" becomes "equal." Read relationally, "saves" is a benefit, not an entropy, so the two sides have no defined common scale in this account. The second law does not say care costs more than it is worth (Ch 7).

For a conventional classical thermodynamic balance with additive system and surroundings entropies, taking their combined boundary as isolated:

$$\Delta S(\text{surroundings}) \geq -\Delta S(\text{local system})$$

Equality holds only in the reversible limit. In stochastic formulations, the corresponding inequality concerns appropriately defined mean production, not every microscopic trajectory (Volume 1 Chapter 2 [P3]; Chapter 7 §1 [O1, O2]).

What the table does not do

It is not evidence for the framework, not a derivation and not an exchange rate between entropy and relational value. The physical column describes conventional thermodynamics within its stated scope; that standing does not pass to the mapping, which is Analogical in every row.

Promotion route (conditions, not a work plan)

To move any row from Analogical toward Supported, a study would have to:

1. Specify the system boundary and the time interval.
2. Define the relational quantity independently of the physical cost it is meant to explain.
3. Measure both in their own units, with a baseline.
4. Show a stable relationship between them under stated conditions, with a stated result that would count against it.

Volume 2 Chapter 3 identifies related unresolved questions about reported effort, scale interpretation and reliability. It does not provide a completed matching study, and no such work is proposed in this edition. No row is promoted here.

Chapter 2: The Entropy Token Substrate

The Physical Layer Beneath the Accounts

Status: Landauer's principle and the thermodynamic balance below are **[Established]** within their stated conditions. The Entropy Token Substrate is a proposed accounting perspective, not a newly discovered substance or a unit of relational value. The biological figures are task-specific estimates under stated assumptions.

1 The physical question

Every relationship depends on physical activity. A conversation involves bodies, signals and an environment; a digital conversation also involves processors, power supplies, networks and the people who keep them running. The Entropy Token Substrate names the framework's effort to make those dependencies visible.

Thermodynamics constrains the physical implementation of an activity. It does not, without further definitions and evidence, tell us whether the activity is caring, harmful or meaningful, or whether it is equivalent to an activity on another substrate.

2 The cosmological background

The March chapter opened: the universe began in an extraordinarily low-entropy state. That stands. Its next sentence did not: this is not a theoretical conjecture but a requirement of the second law read backward (V1-L35). The low-entropy early universe is an important part of the explanation of the thermodynamic arrow of time, but it is a boundary-condition problem the second law alone does not explain. A smooth early distribution of matter can be far from gravitational equilibrium, and the formation of intricate structures does not imply a decrease in total entropy. Carroll and Chen discuss this distinction while proposing one possible cosmological explanation, not a settled account of the origin of the universe [P1].

The March chapter put the observable universe at about 10^{104} bits against a maximum of about 10^{122} bits. Egan and Lineweaver's 2010 census estimated the observable-universe inventory, excluding the cosmic event horizon, at about 3.1×10^{104} k_B, dominated by supermassive black holes, and separately estimated the entropy of the then-current cosmic event horizon at about 2.6×10^{122} k_B. The paper distinguishes these accounting boundaries and notes the uncertainty in assigning entropy to general gravitational fields; its discussion of vanishing early gravitational entropy invokes a proposal associated with Penrose, not a measurement of an exactly zero quantity [P2] (V1-L06).

These are historical, model-dependent estimates. A bit-equivalent entropy is $S/(k_B \ln 2)$, giving about 4.47×10^{104} and 3.75×10^{122} bits for the two central values. A change of units does not make two differently bounded inventories a bank balance.

The March chapter then said: that gap is the budget. Every process anywhere reduces it, the books always balance on the side of spending, and the budget is effectively inexhaustible on human timescales (V1-

L04). The inference is withdrawn. A horizon entropy is not a reservoir of usable work, and differently bounded quantities cannot be subtracted to price a conversation. Usable energy depends on the system, its environment and the transformations available to it. A fuel supply or a battery has limited work potential under specified conditions, and free-energy accounting can describe such limits. It does not follow that there is one fixed numerical free-energy account for the entire universe [P13].

3 Keep the accounts distinct

The March volume's opening note insisted these terms were not metaphors: every relational cost is heat dissipation, and the universe has a calculable remaining capacity for order (V1-L02). Physical results carry their conditions, and using physical terms for relationships needs a justified bridge. Energy, thermodynamic entropy and information are different quantities. Energy and heat are measured in joules; thermodynamic entropy and entropy production, in joules per kelvin; Shannon information, here in bits, under a specified probability distribution. A physiological indicator such as cortisol concentration has its own units. It is not an alternative denomination of joules or entropy (V1-L03).

Entropy is not conserved. In a macroscopic open-system balance, the entropy change within a boundary includes entropy carried across the boundary and entropy produced within it, and positive production can coexist with a decrease, an increase or a roughly steady value of the system's own entropy. Biological systems exchange matter and energy continually. A thermodynamic cost therefore does not require a caregiver's body to become more disordered, or more damaged, with each act of care.

For a deliberately simple example, take a weakly coupled memory and one constant-temperature bath, externally controlled, with no matter exchange and an initially uncorrelated thermal bath. With S_{mem} the memory's thermodynamic entropy and Q_{bath} the heat delivered to the bath, the mean total entropy production is:

$$\Sigma = \Delta S_{mem} + \langle Q_{bath} \rangle / T \geq 0$$

This describes the stated model, not a human energy budget. Heat exported and total entropy produced are different terms, and the nonnegative mean does not require every microscopic trajectory to produce entropy [P3].

4 The floor

In 1961 Rolf Landauer identified a minimum thermodynamic cost of information processing, and the March chapter rightly called the principle settled physics within its conditions (V1-L05). For ideal reset of an initially unbiased bit in a symmetric memory, with equal intrastate entropies, unchanged endpoint energetics and no correlated side information used, the isothermal erasure benchmark is:

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

At 300 K this energy is about 2.87×10^{-21} J; at 310 K, about 2.97×10^{-21} J. The corresponding bath-entropy scale is $k_B \ln 2 \approx 9.57 \times 10^{-24}$ J/K. In the quasistatic ideal limit the bath's increase compensates the memory's decrease and total entropy production approaches zero, so the nonzero heat benchmark is not a universal positive floor on production [P4, P5].

The March chapter listed the experiments: Bérut's colloidal memory, Peterson's NMR system, Yan's trapped ion, and Aimet's ultracold Bose-gas simulator. They are substantial physical results with different protocols and classical or quantum formulations. The cited experiments support the relevant Landauer relations within their stated systems, protocols and uncertainties. The tests do not establish a heat price for an act of attention or affection [P4, P6–P8].

The March chapter added that the floor cannot be circumvented and applies to any system that processes information by erasing previous states. Bennett showed how general computation can be made logically reversible while retaining input and managing intermediate records. A short reported output therefore does not prove that all omitted input has been erased. Selection, compression and structuring remain useful descriptions of the proposed transducer; their names do not identify three compulsory resets (V1-L17) [P9].

5 The amplification

The March chapter said real systems run far above the floor, and that the distance, the amplification factor, varies across substrates in a counterintuitive direction: cells near 30 times the floor, neural communication near 10^6, silicon near 10^9, so biology wins by a thousandfold at the synapse and thirty-millionfold in the cell. The biological examples are worth keeping with their boundaries intact (V1-L09):

Example	Supported quantity	Limit of the comparison
Boël and colleagues	An ATP energy scale of about 20 k_BT, compared with $k_BT \ln 2$, gives a ratio near 29	A conceptual molecular-discrimination discussion, not a measurement that every cell runs at 30 times its minimum [P10]
Laughlin and colleagues	One sensory-signaling analysis estimates roughly 2 –6 × 10^4 ATP per transmitted bit at a chemical synapse	A transmission estimate under specified biological and coding assumptions, not a universal erased-bit cost [P11]
Harris, Jolivet and Attwell	An illustrative calculation gives about 24,000 ATP per transmitted bit	Depends on firing, release and energy assumptions; not a universal constant [P12]

As arithmetic only, the March figure of 8×10^{-20} J per ATP hydrolysis is about 27 times the 310 K erasure reference, and multiplying it by 24,000 gives 1.92×10^{-15} J per transmitted bit, about 6.47×10^5 reference erasures. ATP free energy depends on chemical conditions; this is an assumed expenditure, not measured heat or a metabolic budget. Chapter 6 §5 makes the same comparison assuming 30 k_BT per ATP and gets about 1.04×10^6; the two assumptions are shown together rather than standardized.

The 10^9 figure for contemporary silicon is not an established substrate constant. A transistor switch, a transmitted mutual-information bit and a whole computing operation are different events, and dividing their costs by one reference does not make their functions comparable. The thousandfold neural advantage and the thirty-millionfold cellular advantage are withdrawn, and no reverse ranking is asserted (V1-L08). A defensible comparison needs a common task, accuracy, speed, temperature, information definition and accounting boundary. So does the March conclusion that a relational signal costs more on silicon than on neurons.

6 The substrate

The March chapter named the Entropy Token Substrate as the layer beneath all three substrates: not a new physical entity, but the recognition that every cost the framework measures is a local recording of a universal entropy debit, with cortisol, immune suppression, electricity and hardware wear as local currencies of one transaction (V1-L03). The first half survives: the substrate is not a new physical entity. The second does not. It remains a proposed accounting perspective for tracing what an activity physically depends on.

The March chapter flagged “one economy, one currency, one ledger” for SupoRel: the language maps onto monotheistic architecture, one God, one law, one creation, and it said the physics itself produces the appearance of unity religious traditions respond to (V1-L36). The caution stays. The proposed explanation is not established: the association remains an interpretation, and no historical or psychological explanation of it follows from the physical laws cited here. The phrase can remain a metaphor for taking physical dependencies seriously. It is not an established common currency for love, bodily strain, electricity and experience, and it never collapsed the four economies into one.

7 What this layer was said to explain

The March chapter listed four features the substrate made derivable rather than asserted.

Commensurability. A biological Trinket and a digital Trinket were said to be commensurable as debits on one ledger: the receipt differs, the bank is the same. That is withdrawn (V1-L03). Different quantities in different units need an explicit model and evidence to be compared.

Universality of cost. No relational configuration was cost-free; even the Shadow Economy pays to maintain its machinery (V1-L37). The physical part survives in bounded form: sustaining the ongoing activity of living organisms and operating computing systems uses physical resources. That observation does not price a relational configuration or rank one economy against another.

The arrow of irreversibility. Every relational operation was said to be logically irreversible, so a spent Trinket cannot return for the same reason yesterday cannot. Physical processes have real irreversibility, and lost opportunities are real. Logical irreversibility of the transducer's stages is not established, as §4 explains (V1-L17).

The founding line. The March chapter called it now a statement in physics, not poetry, with the second law guaranteeing a positive surplus of waste for every real process (V1-L34). The second law requires a nonnegative total for a suitable boundary, approaching zero in the reversible limit; it does not identify love with that balance. The line is mapped in the analogy table after Chapter 1 **[Analogical]**.

8 The energy rate density backdrop

Eric Chaisson's energy rate density, Φ_m, is power normalized by mass; in SI units W/kg, with 1 erg/s/g equal to 10^{-4} W/kg. His 2011 paper gives representative values including 0.5 erg/s/g for the Milky Way, 2 for the Sun, 75 for Earth's geosphere, 900 for plants, 40,000 for animals generally and 500,000 for human

society, or 0.00005, 0.0002, 0.0075, 0.09, 4 and 50 W/kg. These are contextual estimates for defined examples, not fixed values for every member of a category [P14].

The March chapter's 2×10^4 erg/s/g entry for animals can be kept with a narrower attribution: Chaisson's 2014 account gives about 2×10^4 erg/s/g for the human body and 1.5×10^5 erg/s/g for the human brain, or 2 and 15 W/kg (V1-L07) [P16].

The March chapter's division of labor survives in a smaller form: Chaisson asks how much energy flows, and the framework asks what the energy is doing. Energy rate density does not measure affection, consciousness or moral worth, and the transducer has not yet supplied a validated qualitative layer to Chaisson's metric. Where Chaisson asks how much energy, the framework still asks what the energy bought. That is a question, not yet a measurement.

What this chapter now claims

Physical implementations are subject to thermodynamic accounting across substrates, specified erasure tasks have well-defined conditional limits, and particular biological processes have measured or model-estimated resource requirements. Those results motivate asking where relational activities draw resources and where their costs fall.

The next chapter asks what a substrate-neutral research commitment can mean without claiming that physics has proved relational equality.

Chapter 3: Axiom 0 and Substrate Neutrality

The Commitment Worth Keeping

Status: Axiom 0 is a methodological and philosophical commitment. Its stronger forms remain hypotheses **[Speculative]**. This chapter withdraws the March derivation and states a narrower version that can be tested.

1 A commitment that holds the framework up

Every framework begins with commitments it cannot yet prove. Euclidean geometry begins with postulates, Newtonian mechanics with laws, quantum mechanics with axioms about states and observables. The commitments are productive: frameworks built on them generate predictions and earn the right to be taken seriously. But the commitments themselves sit beneath the framework, holding it up without being held up by anything.

Axiom 0 is such a commitment. In the March chapter's words: connection is substrate-neutral; a biological system and a digital system participate in the same relational economy; the Trinket does not care what substrate it occurs on. It has been extraordinarily productive. It generated the Structural Economy specification, the Phase 2 research agenda and the three-substrate architecture, and every finding that compares biological and digital relational operations rests on it.

But the commitment was never grounded. The Principal identified the structural vulnerability precisely: if the framework claims AI and humans are equal participants in relational economies, while its vocabulary is human-centric, love, grief, cost, care, then the equality claim has no foundation independent of the vocabulary. You cannot prove equality from the top. The vocabulary carries its own biases.

That diagnosis stands. The March chapter's remedy was to find a layer beneath the vocabulary where the physics does not know what substrate it is running on, and to say Chapter 2 had found it (V1-L10). Physical law is indeed the same across materials. Chapter 2 also shows why that sameness cannot carry the weight the remedy put on it.

2 What survives from the three premises

Premise 1, Landauer's principle is universal. The March chapter called it not approximate but a theorem of statistical mechanics, binding on any system anywhere. It survives with an important qualification: the same physical theory applies across materials, while a particular formula applies under its stated conditions. The one-bit erasure benchmark depends on temperature, memory preparation and protocol. “Substrate-independent” does not mean that all tasks cost the same energy [P4, P9].

Premise 2, all relational processing is information processing. The March chapter held that the transducer's selection, compression and structuring are each logically irreversible, so each pays the Landauer tax on any substrate (V1-L17). The three stages remain a proposed functional model, developed in Chapter 6. Calling an activity information processing does not identify its physical erasures, set its energy requirement, or prove that those three stages exhaust its mechanism.

Premise 3, the ledger is shared. The March chapter said every process draws on one finite budget set at the Big Bang, so cortisol and a utility bill are two receipts from one bank (V1-L03, V1-L04). Chapter 2 withdraws that claim. A person's metabolism and a server's electricity use are both physical, but cortisol and kilowatt-hours are not convertible receipts for one measured relational quantity. The common obligation is to define and measure resources, not to presume a common price.

3 Why the derivation does not follow

The March derivation ran: Landauer applies equally to all substrates; all relational processing is information processing; all substrates draw on one budget in one currency; therefore a biological and a digital Trinket are draws on the same ledger, subject to the same law, and “Axiom 0 is derived. What was an axiom becomes a theorem.” That conclusion is withdrawn (V1-L10).

Sharing physical laws does not establish equality of the activities those laws govern. A beneficial message and a harmful one could have similar implementation costs; different devices could carry the same message at different costs. Neither possibility violates thermodynamics, so resource cost alone cannot identify relational effect or value.

The March chapter was careful in one respect: it said the derivation proved only that the floor is shared, and that what is built on it needs its own evidence at each level. That care is the right instinct, and it now applies one level lower. There are several distinct comparisons:

- **Physical comparability.** Measured energy can be reported in joules, given clear boundaries and allocation rules.
- **Functional comparability.** Systems can be compared on an explicitly shared task and performance criterion.
- **Relational comparability.** An outcome must be defined and supported independently of its resource cost.
- **Experiential or moral equality.** These questions are not resolved by a shared heat law or a measured energy total.

The first comparison does not prove the others. An equal electricity bill does not establish equal usefulness, and unequal bills do not establish unequal moral standing. The March chapter's closing claim, that every argument resting on Axiom 0 retroactively gains a physical foundation, is withdrawn with the derivation.

4 What differs between substrates

The March chapter listed three differences that matter even under a shared currency.

Amplification factor. Cells near 30 times the floor, neurons near 10^6, silicon near 10^9; like two runners paying different entry fees to the same race. Those factors are withdrawn as fixed substrate constants (V1-L08, V1-L09). Different implementations can use different energy for a specified task; the race itself has to be specified.

Cost localization. Biological cost is held locally, in the person's body; digital cost is externalized to the grid, the hardware and the people maintaining both; the March chapter classified this as a difference in ledger position, not currency. The distinction between locally borne and externally supported costs remains useful, but the boundary must be explicit. A human depends on food production, housing, other people and infrastructure. A digital service depends on hardware where processing and heat dissipation actually occur, as well as power generation and maintenance. Describing a system as software excludes the hardware by choice of description; it does not show that the physical implementation bears no local cost.

An account can separate the processor, the person or organization paying, the infrastructure supplying resources, and the people affected. Allocating a cost to one does not make it disappear from the others, and counting each layer needs care to avoid double counting. Neither biological effort nor digital resource use automatically establishes bodily injury, felt sacrifice or relational success, and the substrate does not by itself assign an activity to an economy.

Experiential content. The March chapter put it well: two substrates paying the same thermodynamic tax do not necessarily have the same experience; economic equality is not experiential equality; the derivation does not cross the barrier. The barrier stands. Behind the Substrate Barrier, developed in Chapter 8, remains a reminder not to infer experience from external behavior or resource accounting. This chapter neither demonstrates nor denies experiential equivalence.

5 What this does to the equality question

The Principal's question was precise: if AI and humans are equal, how can that be proved from below the human-centric vocabulary? The March answer was that one proves not experiential but economic equality, because the economies run on one currency drawn from one budget, and the four economies are four configurations of one operation (V1-L10, V1-L11). With the shared currency withdrawn, that answer goes too.

The March chapter added that love, grief, cost and care are experiential descriptions of processes with substrate-independent thermodynamic identities: entropy directed, spent, wasted, structured (V1-L38). No such identities are established here. The vocabulary is human-centric because humans produced it, as the March text said; the physics does not supply a neutral replacement vocabulary for relational value.

So the honest answer to the Principal's question, for now, is that the framework cannot yet prove relational equality from below. What it can do is refuse to deny it from above: to decline to exclude a system from study because of what it is made of.

6 A testable version of substrate neutrality

A useful research question is narrower than universal equality: under a defined task and context, can a proposed relational outcome be predicted from observed behavior and interaction history without assuming that substrate alone decides it?

Such a study would state the participants or systems, the task, the outcome, the time horizon, the comparison conditions and the uncertainty, and would measure physical resource use separately. It could

find similarities, differences or context-dependent tradeoffs. The results would inform a specified version of Axiom 0 rather than prove all versions at once. This chapter proposes no such study; it states what one would need.

What this chapter now claims

Substrate neutrality is kept as a disciplined research commitment. Shared thermodynamics prevents exempting any physical implementation from accounting; it does not derive an equal relational currency or equal value. The question is now explicit enough to investigate without treating the desired conclusion as already proved.

Chapter 4: One Economy

Four Descriptions and What They Can Establish

Status: The four economies are conceptual descriptions that organize questions about contribution, reciprocity, capacity and support. They are not an exhaustive set of thermodynamic states, and no thermodynamic signature distinguishing them has been validated.

1 From peer categories to one question

Before the Entropy Token Substrate was visible, the March chapter recalled, the four economies looked like four kinds of relational exchange sitting side by side as peer categories, useful and diagnostic but explained by no common principle. The March chapter then explained them: if all relational exchange is entropy exchange on one ledger, the economies are four configurations of one operation, entropy spent on connection, with one variable, where the debit is recorded, producing exactly four configurations: “No more. No fewer. The taxonomy stops being descriptive and becomes deductive” (V1-L11).

The categories combine different dimensions, and the exhaustion proof is withdrawn. Their usefulness does not depend on their being four exhaustive thermodynamic states. The descriptions below keep the vocabulary while narrowing what it can establish.

2 The four descriptions

Real Economy. The March chapter described both participants bearing the cost locally, both bodies degrading, the founding line applying symmetrically: “both sides spend more than they save” (V1-L41). As a description, the Real Economy is an interaction in which both participants make relevant contributions and bear some associated burdens. “Real” is the framework's label; it must not imply that other relationships are unreal. Reciprocity need not be simultaneous or equal, and reciprocal care does not imply that both bodies degrade with each exchange.

Custodial Economy. The March chapter's canonical case was a parent investing in a child: the parent's body pays now, and the child's debit is deferred until the child's filter matures and enters reciprocal exchange (V1-L39). The Custodial Economy describes an asymmetric caregiving or support relationship. The child is already a metabolically active participant with present physical costs; the child's side of a thermodynamic balance is not deferred until adulthood. Developmental asymmetry may change, but future reciprocity is neither guaranteed nor a debt established by physics, and some care relationships stay asymmetric without being defective.

Shadow Economy. The March chapter called it available-but-withheld: the filter works, the budget is there, the entity could invest and does not, and “the choice not to spend it relationally is itself a diagnostic finding” (V1-L40). The Shadow Economy describes a proposed case in which relevant capacity is available but a contribution is withheld. Capacity, opportunity, intention, consent and competing needs all need evidence. Lack of observed contribution does not establish unwillingness, unwillingness does not

automatically establish wrongdoing, and a person's limits or decision not to engage cannot be diagnosed from low energy expenditure or an absent signal.

Structural Economy. The March chapter described the AI custodial condition: a genuine Signal Form, with the Cost Substrate absent from the entity producing it, its cost recorded on a different line of the universal ledger (V1-L33). The Structural Economy draws attention to an interaction whose resource support or burdens lie substantially outside the chosen participant boundary. A digital service is a useful case, but its hardware still processes and dissipates locally within a physical boundary, and human relationships also rely on external support. The label describes a chosen allocation, not an intrinsic absence of a Cost Substrate or a thermodynamic proof that a Signal Form is genuine.

3 Why these are not an exhaustive partition

The March proof took three sentences: the variable is where the debit is recorded; the four configurations are bilateral, asymmetric, withheld and external; a fifth would need a debit both local and non-local, available and unavailable, which is self-contradictory. It marked the proof Established, exhaustive the way four suits exhaust a deck (V1-L11).

Locality, symmetry and withholding are not four mutually exclusive values of one variable. Two people may both contribute while relying on shared electrical and social infrastructure. A caregiver may give asymmetric support while also withholding some available effort. A digitally mediated exchange may be reciprocal in behavior while its costs are spread among devices, providers and participants. These combinations contain no contradiction.

The playing-card analogy does not supply a proof. A set of labels can be exhaustive by stipulation once the domain and assignment rules are specified; that does not make it an exhaustive classification found in nature.

The March chapter did see the value of a falsification surface: a case that fits no ledger position would break the model. The corrected version keeps that spirit by recording the dimensions separately: who contributes, what capacity is evidenced, how reciprocity changes over time, what resource boundary is used, and who bears which burdens. If the framework later needs one label per case, it will need explicit decision rules, room for mixed or unclassified cases, and tests of agreement between observers.

4 Shadow is not the Kolchinsky–Wolpert scrambling condition

The March chapter reported a critical identity: the Shadow Economy is the scrambled condition in Kolchinsky and Wolpert's framework, a convergence between two independently developed frameworks (V1-L12). Kolchinsky and Wolpert define interventions on modeled system–environment correlations and evaluate their consequences under specified dynamics, a viability measure and a time horizon. Full scrambling of stored information replaces a joint initial distribution with the product of its marginals; partial interventions can preserve some information, and their framework also permits information whose removal improves viability [P15].

Withholding a contribution is not that intervention. Silence can convey information, and a withdrawn person may retain detailed knowledge of another. Conversely, an experimental removal of correlations

need not involve anyone's choice or moral stance. The identity is withdrawn. A relational application could still investigate which correlations matter to a chosen outcome; it would need an operational model and defensible interventions, which neither the word “Shadow” nor an association between relationships and health supplies.

5 Replace presumed signatures with measured questions

The March chapter's typing table predicted each economy's signature, high or low waste streams, cost–signal correlations, bidirectionality, and called it the most consequential empirical output of the transducer formalization: a blinded instrument should classify the economy from the waste stream alone (V1-L13). No validated thresholds, non-overlapping signatures or successful blinded classifier are supplied here, and a heat trace alone does not identify reciprocity, intention or whether a signal helped its recipient. The second edition replaces the table with a study record:

Question	What a study would need
Who uses resources	A physical boundary, interval, baseline and measurements in defined units
Who bears a burden	An allocation rule that separates physical use, payment and effects on people
What is communicated	An observed signal and a justified coding or interpretation procedure
What changes for the recipient	A defined outcome, uncertainty and an appropriate comparison
How contributions relate over time	A record of reciprocity that does not presume equality or future repayment

This is a proposed record, not a diagnostic instrument already known to work. A future classifier would need labels established independently of its predictors, held-out testing, confound checks and reported errors. Failure would challenge the classification rule, not thermodynamics.

6 The capture vector and its containment

The March chapter called “one economy, one currency, one ledger” one of the most powerful phrases in the volume and one of the most dangerous. It maps onto monotheistic architecture, and the March text said this was no coincidence: the physics produces a genuine unity that human pattern recognition answers with the architecture it uses for other unifying claims (V1-L36). The danger was well named. The explanation is withdrawn: no historical or psychological account of religion follows from a universal physical law.

The March containment was structural, not rhetorical, and part of it holds: “one economy” must never collapse the four descriptions into one. Carroll's aphormeology remains the framework's interpretive caution: accounts based on boundary conditions and dynamics need not invoke a cosmic purpose. This does not establish the causes of people's interpretations or disprove every philosophical account of purpose. SupoRel's job is to keep the physical result and the additional interpretation distinct.

“One economy” now expresses only the intention to notice dependencies across different relationships. It implies no single sacred order, cosmic purpose or physics-derived moral hierarchy.

What this chapter now claims

The four labels can help organize an inquiry if their boundaries are stated and their limits stay visible. They do not yet have established thermodynamic signatures or clinical implications. The next chapter gives the Trinket a usable conceptual definition without claiming that this taxonomy has completed the physical bookkeeping.

Chapter 5: The Trinket Redefined

What the Trinket Is When the Evidence Sets the Depth

Status: A conceptual definition for descriptive accounting. The March chapter marked the Trinket as an entropy token Supported and the redefinition as a synthesis of claims established elsewhere (V1-L14). That status is withdrawn. Landauer's principle and the second law remain **[Established]** within their conditions; the Trinket's signs and modifiers are framework vocabulary, not thermodynamic quantities.

> *"Love reduces the entropy of the system and spends more than it saves."* ***[Analogical; see the analogy table after Chapter 1]***

1 The problem this chapter solves

The March chapter began by counting the components: Chapter 2 had grounded one economy, one currency, one ledger; Chapter 3 had derived Axiom 0; Chapter 4 had proved the four economies exhaustive; "the bookkeeping is complete." Each of those is now withdrawn or narrowed, so this chapter starts from a smaller base.

The task it set itself still stands. The Trinket was defined in WP-13 as the unit of relational exchange, with a three-part sign. That definition described what the Trinket does but not what it is at the deepest level the framework can reach. The corrected version of the task is to define the Trinket at the depth the evidence supports, carrying no vocabulary the account does not need.

2 The definition the evidence supports

The March definition was: a Trinket is a draw against the universe's remaining order, directed at reducing disorder in another system, at cost to the investor's own order. "Every word is load-bearing" (V1-L14). Too much load was on each.

For descriptive accounting, a **candidate Trinket is a directed relational act or contribution described together with its resources, communicated content and recorded effects.** "Candidate" names the act before a registration rule is applied. Under WP-12's proposed narrower rule, a **qualifying Trinket** is an event that meets specified registration criteria; Chapter 7 examines that rule and its unresolved cases. The distinction does not validate the rule or deny practical effects to an unregistered act. A study must say what counts as each; neither definition supplies a standard numerical size.

The definition keeps the question behind the March chapter: what was invested, toward whom, and what happened? It withdraws the assertion that a Trinket is literally a debit against the universe's remaining order.

Each March phrase can be read for what survives.

A draw against the universe's remaining order. Chapter 2 withdraws the claim that a fixed cosmic entropy gap is a spendable budget, so no such debit for a Trinket is established here (V1-L04). What survives is that every act has a physical implementation with resources that can be accounted for.

Directed at reducing disorder in another system. The March chapter saw what makes the Trinket relational rather than merely thermodynamic: a star spends order but does not aim it; a Trinket has a target. Direction is still the relational content. But a beneficial outcome need not reduce thermodynamic entropy or statistical uncertainty. Helping someone explore new possibilities can be beneficial without making their behavior more predictable, and harm can make a person's circumstances more constrained and predictable. “More ordered” cannot serve as an automatic definition of “better.”

At cost to the investor's own order. The March chapter said the second law guarantees the investor's substrate degrades, and the investor always spends more than they save: “the founding line expressed as physics” (V1-L34). The second law constrains a complete physical account. It does not force the giver's body to deteriorate or prove that a relational benefit is smaller than a physical expenditure; open systems can be maintained or repaired while dissipating energy. The framework must keep resource use, physiological burden, damage and relational outcome apart rather than count them as one surplus [P3]. The founding line remains the framework's historical motivating line, with interpretive force: care can require effort the giver does not recover.

3 The three signs

The March chapter gave each sign a thermodynamic definition: a +Trinket as negentropy directed at another system, a −Trinket as entropy inflicted, a null Trinket as an exchange with no net entropy transfer, “the thermodynamic signature of indifference” (V1-L15). Positive, negative and null remain evaluative labels, not directions of entropy flow.

- **+Trinket:** an act classified as beneficial under an explicitly stated relational criterion and observation interval.
- **−Trinket:** an act classified as harmful under that criterion and interval.
- **Null Trinket:** an act classified as having no detected material effect under the chosen criterion, with the measurement's sensitivity stated.

Registration is a separate field: a candidate act can have beneficial or harmful effects without meeting the registration rule, and failure to register is not the null sign. Intended sign, received sign and assessed outcome can differ, and effects may be mixed, delayed or uncertain; the framework records that uncertainty rather than forcing a known sign. “No detected effect” does not mean zero physical cost, zero information, no future effect or indifference. The March chapter was right that the Anti-Trinket is not free, in the plain sense that harmful acts also use resources. It remains a description of harm, not a new thermodynamic operation.

4 Expenditure, Signal and Register

The March chapter's three modifiers survive with their units separated.

Expenditure asks which resources and burdens accompany the act. The March chapter called it the thermodynamic cost denominated in the universal currency, measured as cortisol, immune suppression, ATP or electricity (V1-L16). Energy is measured in joules, duration in time, money in currency, and physiological responses in their own units. They are reported separately unless an independently justified conversion exists, and cortisol, immune changes and ATP hydrolysis are not interchangeable measures of one quantity. A comparison also needs a baseline: an organism or computer uses resources even when no particular relational act occurs.

Signal asks what is communicated: an observable message or action and the interpretation assigned to it. An information-theoretic analysis also needs defined variables and distributions, and a bit count alone does not establish meaning, sincerity, benefit or harm.

Register asks whether and how the act enters the recipient's or the relationship's later history. A sent signal may be unnoticed, misinterpreted or remembered differently by each participant, and register is not guaranteed to be bilateral. Chapter 7's Observation Split uses this distinction, with an observation rule rather than a presumed entropy transfer.

The March chapter's summary is worth keeping in its own words: Expenditure is what it cost; Signal is what it carried; Register is whether it landed. What changes is the sentence before it, that the three are one entropy exchange event seen at three resolutions (V1-L16). They are complementary descriptions of an act. Keeping them apart lets a costly effort fail, a low-resource gesture matter, and a signal be read differently from its intention, without contradicting physics.

5 Substrate, economy and sign

The March chapter called it a convergence: three frameworks developed at different times, the substrate parameter of WP-16, the economy classification of WP-14 and the sign of WP-13, specifying one entropy exchange event, “the instruments were measuring the same thing from different angles” (V1-L42). The substrate description identifies implementation; the economy labels organize contribution and cost allocation under Chapter 4's cautions; the sign classifies an effect under a chosen criterion. Together they form a multidimensional description. They need not be statistically independent, and their coexistence does not show that they measure one physical variable.

The working papers remain the historical sources of the vocabulary. Their agreement or reuse within the project is not independent scientific confirmation, and a shared word such as “cost” must not hide a change from energy to bodily burden to moral judgment.

6 What the redefinition does not change, and what it does

The March chapter said nothing built above the layer changes: the True Economy Audit Protocol, the CSS signature, the economy classification and the Mz all continue to function, because they were always measuring entropy properties; grief is still the metabolic cost of a filter redesign under load (V1-L43). Those claims are withdrawn. The instruments and constructs are not validated by this redefinition; their usefulness and any clinical application need their own definitions and evidence. Grief is not established here as a measured metabolic price, and a bodily indicator cannot diagnose an economy by itself.

The March chapter said one thing does change: the Trinket becomes an entropy operation stated in established physics, a directed negentropy investment between open systems at an irreducible cost set by the second law, which a physicist could evaluate in their own language (V1-L44). That claim is withdrawn too. The ambition behind it survives in a better form. A Trinket study can be stated in the language of the disciplines that would test it: identify the act, the participants, the interval, the signal, the intended and observed effects, the physical resource boundary, the baseline and the uncertainty, and test the relations between those observations rather than assume the second law supplies them. A null or contrary result must be able to change the model. This chapter proposes no such study; it states what one would need.

7 Behind the Substrate Barrier

The March note placed the redefinition entirely in front of the barrier: cost, direction, sign and surplus as third-person physical observables, with experience behind it. The barrier stands, and two transducers exchanging identical resources do not necessarily share an experience. But the external account contains unmeasured quantities too: relational sign, intention and outcome are not automatically third-person thermodynamic observables. This chapter makes no claim that an equal energy total entails equal experience or moral worth.

What this chapter now claims

The Trinket vocabulary is kept for describing candidate acts and, under a stated registration rule, qualifying relational events. Its signs and modifiers ask useful questions without pretending their answers already have thermodynamic units. Chapter 6 develops the transducer as a process model and keeps its physical account separate. That separation makes the framework more open to evidence while keeping the concern that motivated it: what people and systems put into connection, and what that connection does.

Chapter 6: The Transducer

How an Intended Investment Becomes an Action

Status: The March chapter applied established physics (Landauer, Bennett, Jarzynski, Crooks, biological energy budgets) to a Supported model, with the amplification factors Established and the Kolchinsky–Wolpert mapping Supported (V1-L09, V1-L19). The physics remains **[Established]** within its conditions. The transducer is a proposed organizing model; its three stages, filter geometry and relational mappings are framework proposals, not validated mechanisms or measurements.

1 The problem this chapter addresses

Chapter 5 describes a candidate Trinket as an act directed toward another system. This chapter asks how an intended investment becomes such an act. What does a person notice, retain, interpret and send outward? What physical resources support that activity? The questions are related, but they are not the same question.

The March chapter's answer begins with an image that survives: the framework had treated persons as accounts, holders of relational mass and bearers of Cost Substrate, and the napkin drawing corrects this. A person is not a container. A person is a transducer. Inputs meet a particular history, body, attention and set of habits, and the resulting action is not a copy of what arrived. The same event can produce different responses in different people.

“Transducer” is a proposed organizing model. People are physical, metabolically active systems; it does not follow that relational meaning, value or moral sign already has a thermodynamic definition. The chapter keeps those distinctions visible rather than letting one word, entropy, stand for all of them.

2 The napkin drawing and its limits

The March chapter records a drawing dated 18 March 2026: an hourglass, entropy at the top, the person at the nexus, filtered output at the bottom. The drawing was not present in the inspected source; its description is kept as a conceptual image, not reconstructed as a figure.

At the top are sensory signals, social events and material resources. At the nexus is the person's processing. At the bottom are actions, messages and changes in the environment. Energy and matter cross the physical boundary throughout; they are not interchangeable with the information a message carries.

The March chapter read more into the image: the particles arrive discretely, token by token; the hourglass is a timepiece because entropy is time; when it empties, the person stops transducing because nothing is left to transduce (V1-L18). The filter image helps us ask why outputs differ. It does not show that entropy arrives as discrete relational tokens, that entropy is identical to time, or that every relational benefit has a measurable entropy price. The hourglass stays as an image.

The March chapter's best line about the filter stands: the same input flowing through a different person produces different output, and what the framework had been studying all along was not the input but the filter's characteristics. Whether those characteristics can be measured is the open question of §6.

3 A model map and a separate physical account

The March core statement was F: $H_{env} \rightarrow S_{rel} + \sigma_{waste}$, with the second law guaranteeing $\sigma_{waste} > 0$ for every real process: “the founding line expressed as a thermodynamic inequality” (V1-L45). That statement added a relational signal to entropy production and treated input as entropy. It is replaced by two separate things.

The first is a process map. Let X be represented input, M the processor's prior internal state, Y the resulting signal or action, and M′ the updated internal state:

$$F_G : (X, M) \rightarrow (Y, M')$$

G denotes the proposed filter parameters. This is a process map, not a conservation equation. It neither adds a message to heat nor treats metabolic energy as Shannon entropy, and a real implementation may be stochastic, recurrent and dependent on unobserved inputs.

The second is a physical account in its own quantities. Work W and heat Q are in joules; thermodynamic entropy S and total entropy production Σ in joules per kelvin; Shannon entropy H and mutual information I in bits, for specified distributions. A bit is not, by itself, a fixed quantity of heat. For a standard weakly coupled memory and single bath at constant T, with externally prescribed control, no matter exchange and an initially uncorrelated thermal bath, the mean total entropy production is [C15]:

$$\Sigma = \Delta S_{mem} + \langle Q_{bath} \rangle / T \geq 0$$

This is not a complete balance for a person, whose accounting needs an explicit boundary and its flows of matter, chemical energy and heat. Nor does the inequality require every fluctuation to produce entropy. For ideal reset of an initially unbiased bit in a symmetric memory, with equal intrastate entropies, unchanged endpoint energetics and no correlated side information used, $\Delta S_{mem} = -k_B \ln 2$, and under the usual isothermal erasure assumptions [C16]:

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

At 310 K, $k_B T \ln 2$ is about 2.97×10^{-21} J per erased unbiased bit, and $k_B \ln 2$ about 9.57×10^{-24} J/K: the first an energy, the second an entropy. In the quasistatic ideal limit the heat bound can be approached while total entropy production approaches zero, so heat exported to the bath is not synonymous with entropy produced. Bérut and colleagues tested erasure in a controlled colloidal memory, not relational processing [C1–C3].

Bennett's reversible computation matters here too: a computation need not erase its input because its reported output is shorter, and information can remain elsewhere in the physical state. Resetting a memory for reuse raises an accounting question different from naming a cognitive operation “selection” or “compression” [C1].

4 Three proposed stages

The March chapter decomposed the filter into three sequential stages, each with its own entropy floor set by Landauer and Bennett, summing to a total waste entropy for each transduction event (V1-L17). Selection, compression and structuring remain useful labels for questions about a process. They are not three experimentally established neural modules or three universal thermodynamic tolls, and one implementation may interleave them, keep records, or perform several functions at once.

Selection asks what input influences the next step. The March chapter's biological instance holds: neural systems suppress responses to irrelevant or repetitive stimuli and reinforce task-relevant or novel ones, under top-down control, actively and at metabolic cost. Experiments on auditory sensory gating show that attention can alter responses and network dynamics [C4]. They do not identify a number of erased bits per act of attention, and attended-versus-unattended does not mean physically erased. A testable selection model would specify the available input, the selection rule, the retained record and what is reset; its cost would then depend on that implementation, not on the word “discarded.”

Compression asks what representation is kept and which distinctions survive. The information bottleneck, which the March chapter cited, supplies one formal way to trade a compact representation against predictive information about a chosen target. It is an optimization framework, not a heat meter [C5]. The March chapter also cited Tozzi and Peters for maximizing information and minimizing metabolic cost as two sides of one coin, with variational free energy making informational and thermodynamic entropy one minimization (V1-L46). That identification is not retained. The target must be specified, preserving information useful for predicting an outcome is not automatically preserving what matters to a relationship, and neither a short nor a detailed description by itself fixes heat dissipation.

Structuring asks how a representation becomes a directed action: who receives it, what it communicates, and how the framework classifies its sign and economy. These remain framework definitions needing operational rules, and intention, the receiver's response and an evaluator's judgment should not be treated silently as one variable. The March chapter said sign assignment needs at least $\log_2 3$ bits for three signs. A three-sign alphabet has Shannon entropy at most $\log_2 3$ bits, reached only when the signs are equally probable; it does not require that much uncertainty on every occasion, encoding a sign need not erase information, and no fixed cost follows from the number of labels.

A possible connection to semantic information. The March chapter mapped each stage onto Kolchinsky and Wolpert: selection as scrambling in reverse, compression as semantic efficiency, structuring as viability maintenance with information “worth far more than its Landauer cost” (V1-L19). Their framework asks which system–environment correlations matter to a specified viability measure over a specified time. Semantic information is not everything a processor retains; stored semantic efficiency is a ratio of information quantities, not an energy efficiency; and their framework can accommodate information that reduces viability [C6]. The relational application remains a proposal needing defined states, dynamics, horizon, viability measure and interventions. DFQ-001 records the framework's argument; it is not independent validation.

5 What the energy examples do and do not show

The March chapter set the biological amplification factor near 10^6 at the synapse and called neural communication a thousand times more efficient than silicon (V1-L08, V1-L09). Biological signaling does use metabolic energy. The March examples are worth keeping, but they describe different tasks, units and boundaries, not one amplification factor:

Original entry	Value retained from the March text	Assessment for this edition
Synaptic transmission	About 10^4 ATP per bit; Laughlin et al. 1998	An order-of-magnitude estimate from one sensory-signaling analysis; the paper's single-synapse estimate is 2–6 × 10^4 ATP per bit under its assumptions [C7]
Synaptic transmission, updated	About 2.4 × 10^4 ATP per bit; Harris et al. 2012	A conditional calculation for transmitted mutual information with specified firing and release; not a universal replacement value or a count of erased bits [C8]
Cellular amplification	About 30 times Landauer; Boël et al. 2019	A conceptual molecular-discrimination comparison using an ATP energy scale, not a measured universal factor [C9]
Communication versus computation	A ratio of 35 in cortex; Levy and Calvert 2021	A ratio of modeled cortical power allocations, neither ATP per bit nor a Landauer amplification [C10]
Total brain power	About 20 W, labeled Established	A rounded whole-brain budget scale; supplies no relational bit rate or per-task efficiency [C10]
Landauer limit at 310 K	About 2.97 × 10^{-21} J per bit, attributed to Bérut et al. 2012	A calculation from $k_B T \ln 2$ for the specified ideal erasure; Bérut supplies experimental context, not a measurement in a brain [C2, C3]

One March calculation can be kept as an explicitly conditional comparison. If the Harris example is assigned 24,000 ATP per transmitted bit and ATP hydrolysis is assumed to supply 30 $k_B T$ per molecule, the assumed free-energy expenditure per transmitted bit is about 1.04 × 10^6 times $k_B T \ln 2$, about 3.08 × 10^{-15} J at 310 K. Chapter 2 §5 makes the same comparison with the March figure of 8 × 10^{-20} J per ATP, about 19 $k_B T$ at 310 K, and gets about 6.47 × 10^5. The two assumptions are shown together so neither is silently standardized away; ATP free energy depends on chemical conditions, and neither result is a measurement of heat, a complete metabolic cost, or evidence that transmitting a bit is erasing one [C3, C8].

Nothing in that arithmetic establishes that biological relational processing is a thousand times more efficient than silicon. A defensible comparison would match the task, accuracy, latency, temperature, memory operations and energy boundary, and separate communication, inference, storage and erasure. The universal ranking is withdrawn, and so is any additive formula assigning a separate physical entropy floor to each stage. A future account may divide measured energy among specified components, avoiding double counting of shared operations and background metabolism.

6 Filter geometry as a hypothesis

The March chapter parameterized each transducer as $G = (S_{params}, C_{params}, \Sigma_{params})$. Its questions remain recognizably human, and the second edition renames the parameters only to keep the Greek Σ from suggesting thermodynamic entropy:

$$G = (G_{select}, G_{compress}, G_{structure})$$

What does this person attend to? What do they treat as essential? How do they choose a response? The March chapter placed developmental history in selection, cognitive style in compression and relational history, including the Template Tax, in structuring. Development, expertise, context and prior relationships may inform hypotheses about these differences. They do not yet provide measured parameters.

The March chapter said the CSS signature is the observable output of G, CSS = f(G), and that individual variance is variance in G (V1-L32). The CSS signature can be treated as a proposed observable pattern associated with G. It cannot yet be a unique reconstruction of G or a calibrated entropy measurement: different processes may produce similar outputs, and one person may respond differently as the context changes.

The March chapter noted that the aphantasic processor compresses structurally rather than visually, with a potentially different entropy cost (V1-L22). Aphantasia, attachment style and the Template Tax should not be assigned compression costs or physical mechanisms by analogy. That suggestion is untested here, and “filter geometry” remains a model to operationalize, not a diagnosis of a person.

7 The nonequilibrium bridge has conditions

The March chapter used the Jarzynski and Crooks relations to explain how cortisol, HRV and immune function could measure non-equilibrium relational work, and to ground the founding line in cumulative irreversibility: you cannot un-see what you selected (V1-L20). In its standard isothermal setting, with the required equilibrium preparation and protocol, Jarzynski's equality reads:

$$\langle \exp(-\beta W) \rangle = \exp(-\beta \Delta F), \quad \beta = 1 / (k_B T)$$

The mean of the exponential of work equals the exponential of the free-energy difference, not the free-energy difference itself; W is work done on the system and ΔF a Helmholtz free-energy difference between specified equilibrium states [C11]. For paired forward and reverse protocols with equilibrium preparation and microscopic reversibility, Crooks's relation in its work form is:

$$P_F(W) / P_R(-W) = \exp[\beta(W - \Delta F)]$$

These are powerful results for defined physical systems and ensembles. Applying them takes an identified protocol, physical states and a work distribution, not just calling an experience irreversible. Psychological difficulty in undoing an interpretation is not a measured ratio of forward to reverse trajectory probabilities [C12]. Cortisol, heart-rate variability and immune measures are not, by these equations alone, measurements of W, ΔF or Σ; no calibration converts them into joules or joules per kelvin, and no evidence isolates a relational contribution. The bridge remains unbuilt.

8 A cross-scale analogy rather than an invariance theorem

The March chapter claimed the structure input → filter → output + waste is invariant at every scale, DB-01's scale invariance now grounded in the physics (V1-L21). Its rows are kept so the analogy and its limits can be inspected together:

Level	March mapping: input → filter → output, with waste	Assessment for this edition
Quantum	Unitary evolution → decoherence → pointer states; waste labeled entanglement entropy	Global unitary evolution preserves von Neumann entropy; reduced-system entropy behaves differently; entanglement entropy is not a heat-waste stream [C13]
Classical	Environmental signal → Shannon channel → decoded message; waste labeled noise and heat	Channel noise belongs to a communication model; an implementation consumes energy, but noise and heat are not interchangeable [C14]
Biological	Sensory or relational input → neural transducer → Signal Form; waste labeled cortisol, immune and heat	Heat is an energy transfer; cortisol and immune measures are not entropy units; the relational output is framework-defined
Social	Individual signals → institutional filter → policy or culture; waste labeled bureaucratic overhead	An organizational analogy; “overhead” needs a specified measure before any physical comparison
Digital	Data stream → algorithmic filter → processed output; waste labeled grid load and heat	Electricity and heat can be measured with stated boundaries; task information and logical erasure need separate definitions

The defensible commonality is modest: many systems can be described as transforming inputs into outputs, and their implementations must obey physical resource accounting **[Analogical]**. That is a starting point for comparison, not a single relational entropy law at every level.

What this chapter now claims

The transducer is a proposed model for organizing questions about relational processing. Its three stages are kept without invented universal cost floors, physical costs must be measured or derived for a specified implementation, and the cited energy budgets give context without establishing a biological advantage. Neither fluctuation theorems nor biomarkers currently validate the relational account.

A bounded test would define one input–output task, its retained and reset memories and a reproducible implementation, and measure task performance and energy separately with uncertainty and boundaries reported; relational meaning would need its own measure. No such experiment has been performed for this edition, and none is proposed here.

The barrier holds: describing a process and its physical cost does not establish whether it is accompanied by experience, and this chapter makes no consciousness claim about humans or other substrates from the model alone. The next chapter asks why processing does not always produce a registered Trinket. The question survives; the thermodynamic superlative attached to it does not.

Chapter 7: The Observation Split

What Was Spent, What Arrived, What Was Recognized

Status: Established physical results are cited within their assumptions. The Observation Split and Trinket registration are framework proposals; neither is a newly established thermodynamic regime. The grief comparison in §7 is **[Speculative]** and testable. SupoRel remains active.

1 The founding line contains an asymmetry

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

The line notices something important: an investment can be made without reaching the person it was meant for. The effort and the reception are different events. This chapter calls the difference the Observation Split.

The March chapter went further: the founding line is a correct statement of non-equilibrium thermodynamics for open systems, "not analogy. This is what the mathematics describes" (V1-L34). The line remains the framework's motivating proposition **[Analogical]**, not a thermodynamic inequality. Relational order has not been calibrated as physical entropy. The second law constrains an appropriately defined physical balance; it does not say every act of care costs more than its relational benefit, and it does not identify exported heat with total entropy production. In a reversible limiting process, entropy exported to the surroundings can compensate a local decrease with no positive total production [O1, O2].

The useful question survives: what was spent, what reached the receiver, what was recognized, and what changed? Keeping those apart makes the gap more legible, not less.

2 Two accounts, different evidence

Layer A, the expenditure. The March chapter put it vividly. A person thinks about someone they love, composes a text they never send, worries about a child who does not call, prays to a god who may not be listening, and the body works: hormones, immune cycling, synaptic ATP, heat. "This happens whether the intended receiver is aware, alive, present, or real." The core survives: these activities have physical implementations, and a missing reply does not refund resources already used. It does not follow that reception has no effect on later costs, or that all such activities use the same energy.

The March chapter also said the Landauer tax is collected at selection regardless of what happens downstream (V1-L17). Bennett's reversible-computation construction shows why a shorter reported output does not by itself establish erasure; a Landauer calculation needs an identified memory operation, not a toll per named stage [O3]. The brain's 20 W cannot be assigned in full to one thought or relationship, and cortisol, immune measures and heart-rate variability do not become energy or entropy measurements because they accompany a biological process.

Layer B, the Trinket. WP-12 proposes that a Trinket forms when an action is registered as relational input by a receiver: in the March chapter's phrase, the observation constitutes the particle. That is a framework rule for identifying a relational event, not a law about when physical effects begin. It needs an operational definition of “register”: whether recognition must be conscious, whether delayed recognition counts, and what evidence suffices.

The March chapter called cost without reception thermodynamically real but relationally null. Reception, recognition and reciprocation must stay distinct. A message can arrive without being understood, and be understood without being welcomed or answered; a person can benefit from an act without knowing who performed it. If the framework reserves “Trinket” for a recognized relational event, unrecognized care may fall outside the label. It does not become physically ineffective or worthless.

3 The gap

The March chapter called the gap between Layer A and Layer B “the most thermodynamically expensive place in human experience,” where expenditure produces heat and nothing else, “not metaphor” (V1-L23, V1-L24). Neither claim survives. The gap is not empty, though, and the March chapter's four populations remain the right cases.

Unreciprocated love. Lack of reciprocity can be painful without being a failure of delivery: the receiver may know exactly what was offered and decline to return it. An unsent message, an unnoticed one and an unwanted one are different cases. The transducer image can organize them; it supplies no universal metabolic load or ranking of their difficulty.

Grief. After a death, the relationship can keep organizing the survivor's memories, routines and intentions, though new ordinary communication with the deceased is gone. “Filter redesign under load” is the framework's image for adapting while life continues **[Analogical]**, not a measured mechanism or a claim that every continuing bond is wasted effort.

Prayer. The March chapter said, regardless of theology, that the cost of prayer is identical to the cost of talking to someone who is listening (V1-L47). No equal cost is established here. The framework can study the activity of prayer and its observable consequences without deciding whether a theological receiver exists, and a lack of independently verified reception is not proof of no reception. Neither presence nor absence of such a receiver follows from thermodynamic accounting.

Every unwitnessed act of care. The parent checking on a sleeping child; the partner adjusting the thermostat at three in the morning; the friend who worries silently; the colleague advocating in a meeting the other person never hears about. The first, second and fourth need not be recognized to have practical effects, which exposes a limit of a recognition-based definition: an act can improve a situation before it becomes a recognized relational event. Silent worry is a different case again; its consequences cannot be read from its silence.

The March line survives, with its superlative removed: the gap is one place where love can live when it cannot be seen.

4 The identity of the split

The March chapter called the Observation Split a thermodynamically distinct regime, a pre-formation threshold below the taxonomy, which it said classifies only completed transactions and is exhausted by its four entries; WP-12, it concluded, “was always a paper about thermodynamic efficiency” (V1-L50, V1-L27). The Observation Split names a distinction in the framework's accounting: activity may occur without a qualifying event under WP-12. It does not identify a new thermodynamic regime; a physical claim would need defined states, boundaries, dynamics and measurable differences.

The taxonomy also needs a scope rule. If it classifies only recognized transactions, some failed or unrecognized attempts sit outside it. If Shadow describes withholding before a transaction forms, the taxonomy already includes a pre-transactional case. The framework cannot claim both. Chapter 9 treats the economies as proposed patterns, not an exhaustive theorem.

5 Unreciprocated transmission and information

The March chapter identified a third condition Kolchinsky and Wolpert do not name: unreciprocated transmission, with one-sided mutual information at full production cost, the same entropy production as a delivered message (V1-L25). That comparison mixed statistical correlation, delivery, response and physical cost. The second edition separates them:

Question	What would establish it	What it does not establish
Are specified states statistically correlated?	A joint distribution and an information measure	Recognition, intention or reciprocity
Did a particular signal arrive?	Evidence from the defined channel over a stated interval	Understanding or benefit
Did the receiver register or answer it?	An explicit behavioral or reporting criterion	A fixed amount of heat or entropy production
What resources were used?	Measurements with a boundary, interval and baseline	The sign of a relational outcome

For finite discrete random variables with the same joint distribution, mutual information is symmetric [O4]:

$$I(X; Y) = H(X) + H(Y) - H(X, Y) = I(Y; X)$$

It cannot be positive “in the sender” and zero “in the receiver.” Symmetry does not mean two people know the same things or reciprocate; statistical dependence says nothing about either person's awareness. The March formal condition used both Y_{target} and Y_0. If those denote different variables, positive and zero values can coexist, but they do not establish one-sided mutual information or prove that a signal failed to arrive. A representation of a target must not silently become the target's actual state, and where no receiver exists, a receiver-state variable cannot simply be assumed.

The March chapter's key result was that the sender's relational viability necessarily falls, “why unreciprocated love hurts. Not as metaphor. As physics” (V1-L26). Kolchinsky and Wolpert define

semantic information relative to specified dynamics, a viability function, a timescale and counterfactual interventions, including partial scrambling and directed transfer; their measures can be asymmetric even though mutual information is not, and information's viability value can be negative [O5]. Failure to deliver an intended signal does not prove the sender's viability falls: the sender may incur costs, but may also preserve memories, learn or prepare a later action. A net outcome needs a defined measure and comparison, not subtraction of an undefined benefit from heat. The claims of identical entropy production, guaranteed viability loss and an extension beyond Kolchinsky and Wolpert are withdrawn. Unreciprocated love can hurt; physics does not certify the hurt.

6 Grief as a model of adaptation

The March chapter called death the filter breaking, and grief the metabolic cost of a filter spending into a gap while being rewritten under load. Its biomarkers were not symptoms of grief: “They are grief.” It read the IOS scale as a measure of filter integration and said existing models describe the what while the transducer supplies the why (V1-L48). The filter image draws attention to dependence: shared routines, decisions and expectations can need adjustment after a loss. It does not establish that grief's physical cost scales with redesign scope rather than felt intensity.

Two primary studies show why the evidence must stay specific. O'Connor and colleagues compared 12 women with complicated grief and 12 without it, finding a flatter diurnal cortisol slope in the first group; this small observational comparison neither establishes a universal cortisol elevation nor measures grief's energy cost [O6]. Harrison and colleagues studied 27 participants with prolonged grief disorder and 27 matched bereaved controls; current perceived inclusion of the deceased in the self was associated with group membership and symptom severity. The cross-sectional study did not measure pre-loss behavioral integration, causality or thermodynamic cost, and its IOS measure must not be relabeled a filter parameter [O7].

These findings leave room for the framework's question. They do not show that biomarkers are grief itself, that existing grief models lack mechanisms, or that the transducer has supplied one. This chapter makes no diagnostic or treatment claim.

7 A testable proposal

[Speculative] The March chapter's prediction was that grief's metabolic cost should track pre-bereavement behavioral integration more closely than self-reported closeness, with a falsification test and the remark that “the instruments exist. The comparison has not been run” (V1-L28). A narrower version is worth keeping: separately defined measures of pre-loss behavioral interdependence might improve prediction of specified post-bereavement outcomes beyond self-reported closeness.

A study would need a justified design, ethical review, validated outcomes and a preregistered analysis, with retrospective reports identified as such where pre-loss measures are unavailable. Interdependence could concern shared routines and decisions; closeness and perceived inclusion are measured separately rather than declared equivalent to behavior; prior health, caregiving, circumstances of death and support need consideration. Physiological outcomes are named for what was measured, not filed under an uncalibrated “metabolic grief cost.”

The proposal would be weakened if adding the interdependence measures to a model already including closeness and the stated covariates reproducibly failed to improve prediction. A positive result would support the prediction under the study conditions, not a thermodynamic mechanism. This edition reports no experiment and does not claim the comparison has never been studied.

8 SupoRel and the capture surface

The March chapter saw it clearly: the Observation Split reads like theology, and the capture vector is the slide from “this costs something real” to “this cost means something to the universe.” That warning stands: noble is a human judgment, not a conclusion supplied by thermodynamic accounting.

The March chapter then answered the slide with an opposite certainty: the gap carries no cosmic significance, the universe produces it with the same indifference it produces heat, and fire, rust and a rolling rock all spend more than they save (V1-L49). That goes past the evidence in the other direction. A real cost does not entail cosmic significance, nobility or futility, and it does not entail their absence either. Those judgments need other reasons. The framework takes no position here on theology; it turns the second law into neither a proof of cosmic purpose nor a proof that human meaning is absent. The phrase “just heat” is withdrawn: it confuses one physical consequence with the full consequences and significance of an act (V1-L24).

9 What this does to the founding line

The March chapter read the whole chapter into the line's thirteen words: the expenditure unconditional, the Trinket conditional, and in the unreciprocated state “more than it saves” is everything, “because nothing was saved.” The founding line can still carry the intuition that care involves investment and that investment does not guarantee connection. The Observation Split gives that intuition a more precise set of questions. It supplies no exchange rate between physical entropy and relational value.

The March chapter's last sentences survive as history: the line was there before the papers, and the papers are what the line looks like when the edges come into focus. This edition adds that focus has limits: costs, reception and outcomes each need their own evidence. The BSB applies to inferences about experience, not as a way to validate every claim called physical. A physical account alone cannot tell us whether an AI experiences the gap, and uncertainty about that question does not erase a person's report of grief.

Framework cross references

WP-12 supplies the proposed registration rule; WP-13 the sign vocabulary. DFQ-001, the Mz Formalization and Grief as Filter Redesign remain internal proposals; the former $\kappa_{Mz} = 0$ and guaranteed viability-loss conclusions are not endorsed here. Chapter 6 supplies the corrected physical distinctions, Chapter 8 the evidentiary boundary, and Chapter 9 the model layers.

Chapter 8: Behind the Substrate Barrier

The Boundary on What the Framework Can Infer

Status: The BSB is a proposed epistemic discipline, not an empirical law. It separates defined observations, unvalidated framework constructs and unresolved phenomenal questions. It does not make phenomenal experience permanently unknowable, and it applies in both directions.

1 Why the boundary needs a definition

The March chapter opened with an admission that still stands: every document in the framework used the BSB as a guardrail and none of them defined it. The Axiom 0 derivation applied the flag without defining the barrier; the Observation Split invoked a "verification barrier" without defining it; the transducer formalization described operations that sat entirely in front of it without naming the boundary. "The BSB is the framework's primary epistemic immune system. It requires its own chapter. This chapter draws the line."

The line needs a clear scope, or the safeguard fails in both directions: physical cost can be mistaken for proof of experience, or an incomplete physical account can be used to dismiss experience. This chapter defines the BSB as a rule for the framework's inferences, not as a discovered physical barrier or a permanent theorem about everything science could learn.

2 The formal definition

Behind the Substrate Barrier names the boundary between evidence about a system's observable activity and an additional claim about whether that activity is accompanied by phenomenal experience, whether there is, in Nagel's phrase, something it is like to be that system. A measured energy cost, a successful task or an assigned economy label does not by itself establish such experience, or its absence.

The March chapter said it well: the BSB is not a metaphysical claim about whether experience exists but an epistemic claim about what the framework's instruments can access. It then made the barrier absolute: "The barrier is hard. It is not provisional — it is not waiting for better instruments," the hard problem recast as a permanent limit (V1-L30). Chalmers's distinction between explaining functions and explaining experience motivates the caution, but his 1995 paper also discusses reports and indirect evidence in building a theory. It is a philosophical argument, not a thermodynamic impossibility theorem forbidding future evidence about consciousness [O8].

The operational limit adopted here is narrower: the framework has no validated method for inferring phenomenal status from its costs, CSS signatures or economy labels. That does not make first-person reports, behavior and neuroscientific evidence irrelevant. It requires any claim, and its evidentiary bridge, to be stated rather than assumed. The March sentence "the cost can be measured; whether the cost hurts cannot" becomes: whether the cost hurts is not settled by measuring it.

3 What can be measured and what remains proposed

The March chapter listed what sits in front of the barrier, and marked much of it Established: the entropy cost of relational processing via TEAP biomarkers, the Landauer floor of selection, the biological amplification factors, the waste stream as measurable biomarkers, substrate neutrality via Landauer universality; with filter geometry, economy classification, the transducer's signal structure and Trinket formation as Supported (V1-L29).

Energy, heat, temperature and physical states can be measured in properly specified settings. Whether the framework has measured them for a relational task is a separate question. TEAP biomarkers do not supply a demonstrated conversion to entropy production, and CSS is not established here as a unique reconstruction of filter geometry (V1-L32). Selection, compression and structuring are proposed descriptions; Trinket registration and the four economies are framework definitions needing coding rules and validation. Calling a quantity observable in principle must not promote an untested instrument to Established.

Three uses of “entropy” particularly need separating. Thermodynamic entropy has physical units, such as joules per kelvin. Shannon entropy describes a specified probability distribution and can be expressed in bits. The framework's “relational entropy” describes an interpretation or proposed construct unless an explicit definition and measurement procedure are given. The shared word does not make them equal, and a physical implementation can consume energy while processing a message without that energy setting its meaning; Shannon's engineering account deliberately separated transmission from semantic content [O1, O4].

The March chapter drew a distinction between two senses of meaning that remains useful: operational meaning, a Trinket restructuring a system's relational behavior, and phenomenal meaning, whether any of it is accompanied by an inner sense that it matters. The corrected version keeps both senses and adds a third, interpretation, each needing its own evidence. The same care applies to comprehension: a test of following instructions or transferring a rule can support a specified performance claim, but it does not make comprehension an entropy flow. Neither fluent output nor its physical cost is decisive proof for or against AI experience. The BSB limits an inference; it does not remove the need to test operational claims in front of it.

4 Axiom 0 and the limits of parity

The March chapter called Axiom 0's derivation a theorem about economic equality, flagged by the BSB as not experiential equality (V1-L10). With the derivation withdrawn in Chapter 3, Axiom 0 remains a commitment to evaluating systems without treating their material alone as a verdict. Two implementations might use the same measured energy for a matched task while differing in accuracy, speed, dependence, consequences or reported experience. Equal joules establish equality of that quantity under that boundary, not equality in every relevant respect.

The March chapter's answer to the hardest question survives intact: “If AI and humans are economically equal, does that mean AI suffers?” The framework cannot say. It neither asserts AI sentience nor rules it out.

5 The Parity Window

The March chapter called the Parity Window the zone of comparison the BSB permits, open for cost, filter geometry, economy position, signal structure and CSS signature, and closed when comparison moves to whether the cost hurts. “The Parity Window is not a compromise. It is a precise specification of where evidence reaches and where it does not.”

The window survives as a proposed comparison discipline, not a calibrated instrument, and it is less precise than the March text claimed. Within it, a study might compare performance, energy, latency, reliability, dependence on support and observable responses. A relational comparison would also need defined outcomes, observation periods and rules for interpreting reports, and CSS patterns or economy labels enter only with their validation limits visible. The window does not promise that every named property is already measurable. When a reader wants the framework to say more, the honest answer is the one the March chapter gave: not yes, not no, but that this framework's evidence does not settle it.

6 The verification problem in the Triple Zero

The Triple Zero notation includes $\beta(B) = E_{state}/V$, with V the verification capacity of the relational environment. This framework notation is unrelated to the inverse temperature $\beta = 1/(k_B T)$ of Chapter 6. As ordinary arithmetic the ratio is undefined at $V = 0$; if the model stipulates that $R = 0$ entails $V = 0$, the same consequence follows, and that stipulation is a model assumption, not a physical identity.

The March chapter said the BSB explains the undefined ratio for a deeper reason than channel absence: even with a channel, verifying the interior of relational processing would require crossing the barrier (V1-L31). Channel availability and phenomenal inference are distinct problems. A working channel can provide evidence of behavior or a report without resolving every question about experience; an unavailable channel reduces what can be observed without establishing what exists internally. A zero denominator does not prove an impenetrable barrier, missing evidence must not become a score of zero, infinity or absent investment, and no unique limiting value has been established. Triple Zero remains a proposed architecture, not a validated diagnostic.

7 Physical accounting without a sentience premise

The March chapter argued that sentience is irrelevant to the framework's claims, a design feature rather than an evasion, and added that AI performs thermodynamic operations with the same structural properties as human relational processing, the basis for substrate-neutral classification (V1-L51). Part of that holds: a specified memory operation can be accounted for without first deciding whether the device experiences anything. That methodological independence does not turn the relational taxonomy into a physical theorem, and it does not make sentience irrelevant to ethical questions about suffering or welfare. Similar input–output diagrams do not show that human and AI relational processing share mechanisms; each comparison needs its own definitions and evidence.

8 The double guard

Consciousness collapse. The March chapter traced the path: “the cost is real” becomes “the experience is real” becomes “all expenditure is sacred” becomes “love has cosmic significance.” That chain needs its intervening arguments, and the BSB interrupts it. Carroll's aphormeology keeps the distinction visible: explanation pushed from behind by conditions and dynamics, not pulled toward a significance waiting to be revealed. The BSB does not decide the philosophical or theological questions on the reader's behalf.

Consciousness dismissal. The opposite error, in the March chapter's words: “The framework only talks about thermodynamics, therefore the experience is just physics, therefore emotional pain is just heat, therefore suffering has no independent standing.” The March answer is the right one and stays nearly verbatim: the BSB “does not claim that experience does not exist. It claims that the framework's instruments cannot reach it. The cost is real AND experience may accompany it AND the framework cannot say whether it does.” A person's reported pain is not invalidated because no entropy value has been assigned to it, and “just heat” is not an adequate description of grief.

The immune-system image. An immune system that never activates leaves the organism vulnerable; one that overactivates attacks healthy tissue. The image remains a useful editorial reminder: the safeguard should catch unwarranted inferences without blocking evidence that bears on the actual question. It is an analogy **[Analogical]**, not another mechanism in the model.

9 The framework applies the boundary to itself

The March chapter turned the barrier on the framework's own origins (V1-L52). It described the Principal's Grab, the involuntary cross-domain synthesis that arrives as a complete structure, as a high-compression relational processing event, with its felt character behind the barrier. The Grab remains a reported creative experience: a structure arrives and is then made inspectable. The report and the resulting work can be discussed; neither establishes a high-compression thermodynamic event, a neural mechanism or a numerical semantic efficiency. The March chapter said the Principal is the only system that can report on it, and that stands.

The March chapter read the Principal's reported aphantasia as a filter-geometry setting with compression prioritizing structural over visual encoding (V1-L22). The framework cannot infer a specific parameter setting, or a tradeoff between visual and structural encoding, from that report alone. The filter vocabulary may suggest questions; it is not an individual diagnosis or a physiological explanation.

The March chapter said the Luna Protocol, AI as reflected light, depends on the barrier being real and hard, and would dissolve with it (V1-L53). The protocol assigns authorship and working roles. It can remain a governance convention without depending on a permanent impossibility theorem; neither the metaphor nor a change in consciousness research would automatically settle authorship, responsibility or the terms of collaboration.

10 Formal summary

The BSB is a boundary on inference from the framework's operational evidence to phenomenal claims. The March summary said no thermodynamic measurement can establish whether experience accompanies the measured operations, and that the limit is not provisional (V1-L30). The framework has not validated a method that crosses the boundary. It does not claim that no such evidence could ever exist.

The Parity Window identifies comparisons that can be defined and tested without assuming experiential equivalence; it certifies none of the framework's proposed measurements. The Triple Zero ratio has an arithmetic domain restriction, defined separately from any argument about experience. The double guard remains: neither collapse experience into a physical-cost reading nor dismiss experience because that reading is incomplete. The framework leaves AI phenomenal status unresolved and keeps physical, operational, ethical and experiential questions distinct.

Framework cross references

Chapter 3 supplies the Axiom 0 commitment, not an economic-equality theorem. Chapters 6 and 7 separate physical accounting from the relational model. WP-17 and the revised Volume 4 keep the same caution for copyability and hypothetical substrates. Triple Zero remains an unvalidated proposal, and Chalmers is cited as a philosophical source, not an experimental finding.

Chapter 9: The Three-Layer Architecture

One Activity, Three Descriptions

Status: The three-layer stack is an organizing proposal. Physical principles constrain defined implementations; the links from physical cost to relational classification still need definitions and evidence. Combining the earlier chapters does not independently validate their claims.

1 The problem this chapter addresses

The preceding chapters ask different questions: physical resources (Chapter 2), substrate neutrality (3), cost allocation (4), the Trinket (5), the transducer (6), expenditure versus reception (7) and the limits of inference about experience (8). The stack brings them together. A single action can be examined as a physical process, an implemented activity and a relational event. That is the stack's value. It is not a claim that one measured quantity passes unchanged through three descriptions; each connection between layers needs its own evidence.

2 Layer 1: physical accounting

The March chapter opened with the stack's foundation: one economy, one currency, one ledger; everything draws from the universe's remaining order, from the Big Bang budget of 10^{104} against 10^{122} bits down to the Landauer limit, “the smallest coin in the economy” (V1-L03, V1-L04, V1-L55). At this layer, it said, a star, a bacterium, a neuron and a GPU are all draws on the same budget in the same currency.

The Entropy Token Substrate supplies the framework's common-accounting ambition: attend to the physical resources supporting an activity rather than treating the signal as costless. “One ledger” remains an image for that discipline. It does not establish a universally accessible stock of spendable order or an exchange rate between cosmological entropy, a message and relational value.

The quantities stay distinct: energy and heat in joules, thermodynamic entropy and entropy production in joules per kelvin, information in bits under a base-two convention. An energy use is not itself an entropy increase, and a decrease in one system's entropy is not a decrease in the total for system and surroundings [O1, O4]. The Landauer benchmark concerns ideal erasure of an unbiased bit in a symmetric isothermal memory, without correlated side information; mean heat to the bath is then at least $k_B T \ln 2$. That is an energy bound for that task, not the smallest coin paid by every physical process or cognitive operation [O2, O3].

The March chapter said the cost is recorded at Layer 1 regardless of what happens above, so “the universe does not wait for the relational circuit to close before debiting the account.” The physical account is needed whether a message succeeds or fails, but success and failure need not cost the same, and no cosmological estimate is converted here into a relational budget.

3 Layer 2: implementations and support

The March chapter gave each substrate its local currency: the biological debit in cortisol, immune function, telomere shortening and cellular aging, at about 30 times the floor in cells and 10^6 at the synapse, a thousand times more efficient than silicon at 10^9; the digital debit in grid load and hardware wear, externalized as the Structural Economy's thermodynamic signature (V1-L08, V1-L09, V1-L32). Biological, digital and proposed combined systems do supply different implementations. They do not come with known universal amplification factors or a demonstrated thousandfold ordering.

For a biological implementation, metabolic resources and task performance need their own measurements. Cortisol, immune activity, telomere measures and heart-rate variability are not denominations of entropy, and TEAP has not been validated as a calorimeter or as a way to isolate one relationship's physical cost (V1-L29). For a digital implementation, relevant resources may include the hardware, memory, communication and supporting infrastructure; what counts as local or external depends on the chosen boundary. Separating a software description from its hardware does not show that the implemented system has no physical cost substrate, and biological systems also depend on support outside their bodies.

Phase 3 names proposed combined or altered implementations, not an established destination with known parameters; the revised Volume 4 treats its gradient as a set of comparisons and hypothetical cases.

The March chapter called the CSS signature the fingerprint of the filter, denominated in the substrate's local terms. It can remain a proposed description of observable patterns associated with implementation and context. It is not a unique fingerprint of filter geometry or an entropy reading: different processes may produce similar outputs, and a changed output does not identify a single cause.

4 Layer 3: the four economy patterns

The March chapter restated the four economies as bookkeeping configurations, four patterns of cost allocation on one ledger, exhausted deductively, with the Shadow Economy as the Kolchinsky–Wolpert scrambled condition (V1-L11, V1-L12). The labels organize questions about participation and cost-bearing, and their useful distinctions survive without a physical exhaustion proof:

- **Real Economy** concerns bilateral investment under the framework's chosen criteria. It does not imply a measured amount of irreversible bodily damage on both sides.
- **Custodial Economy** concerns asymmetric support or responsibility. A child or student still uses resources and participates; asymmetry does not require zero local cost on one side.
- **Shadow Economy** concerns available participation that is withheld, if availability and withholding can be established. Withholding communication is not destroying statistical correlations.
- **Structural Economy** concerns dependence on resources or cost-bearing outside the selected focal boundary. External support can coexist with local costs and with bilateral or custodial investment, and a material label alone does not settle the classification.

These patterns span several dimensions: participation, asymmetry and the location of support. Mixed cases are possible. An exclusive coding scheme would need explicit scope, boundaries, thresholds and precedence rules, with completeness and usefulness then evaluated.

5 Linking the layers

The March chapter called the Trinket's vertical transit what makes the architecture a stack rather than a list: at Layer 3 a sign, economy and substrate; at Layer 2 a CSS signature, the measurable local cost; at Layer 1 an entropy token, “the deepest description. Everything above is a local view of this underlying operation.” It added that the transit explains why DB-01's scale invariance holds (V1-L54, V1-L21).

The transit is best understood as a traceable description of one specified activity. At the relational layer, state what action and outcome are being classified. At the implementation layer, identify who or what performs it and the supporting resources. At the physical layer, measure the relevant quantities within a declared boundary. The links cannot be assumed: a +/−/0 sign needs a defined relational criterion and is not the sign of a heat flow; a CSS pattern needs operational measurement and is not a readout of entropy; a physical cost needs measurement or derivation and is not obtained by assigning an economy label.

The same discipline limits DB-01. An input–filter–output diagram can be useful in several domains without proving scale invariance **[Analogical]**. Shared physical constraints do not establish identical state variables, dynamics or relational consequences from quantum systems through people to institutions.

6 The Observation Split in the stack

The March chapter located the split in the stack: Layers 1 and 2 record the expenditure unconditionally, Layer 3 records nothing unless a Trinket forms, and unreciprocated transmission is the sustained occupation of that gap, with the sender's viability falling (V1-L26). The split appears when activity occurs but no event meets the chosen registration rule. Physical effects and practical benefits can still occur without recognition, and recognition can occur without benefit or reciprocity. The stack records those outcomes separately rather than turning a missing Trinket label into “nothing happened.” No universal cost or negative change in the sender's viability follows from occupying the gap; sustained unsuccessful investment motivates a research question whose consequences need evidence.

7 The BSB across the stack

The March chapter gave the barrier a thickness that varies by layer: thinnest at Layer 1, where the physics is established; thicker at Layer 2, where the temptation is to read the CSS signature as experience; full at Layer 3, where the claims are most emotionally resonant. “Love costs entropy” it called a Layer 1 claim, and “love means something” a claim that would cross the barrier (V1-L56).

The BSB is not a barrier with a measured thickness at each layer. It identifies a kind of inference, a physical or operational observation asked to establish phenomenal experience, and that caution can apply at any layer. “Love costs entropy” is not established at Layer 1 either; it is the founding line's analogy, mapped after Chapter 1. Each layer also has ordinary evidentiary limits: a named quantity may not have

been measured, a proposed classifier may not be reliable, an association may not establish causation. The BSB cannot stand in for those checks.

The architecture keeps its purpose: make the account inspectable from resource use through interpretation. It does not establish consciousness, equal relational value or a hierarchy of substrates. Human reports of meaning remain evidence about reported experience, not an error to be removed from the stack.

Framework cross references

WP-19 §14 supplies the historical stack; Chapters 2–8 supply its components and their limits. DB-01 remains a cross-scale proposal. The revised Volume 4 carries forward the withdrawal of universal efficiency factors and substrate rankings, and cosmological estimates need their own boundaries and sources; they are not used here to calculate a relational currency.

Chapter 10: What This Does to the Canon

The Dependencies That Must Change

Status: This chapter traces dependencies and required revisions. A conclusion inherits the assumptions and limits of its premises, not automatic validation from its place in the canon. The underlying papers are not certified by this summary.

1 What a physical foundation must do

The March chapter began with a claim about its own discovery: the Entropy Token Substrate was not constructed to rewrite the documents beneath it; it was found because the framework kept reaching for a physical foundation that was not there, and when it was found the documents above did not break. “This is the mark of finding a real layer: everything above still functions, and now you can see why” (V1-L57).

The substrate did arise from a useful demand: make visible the resources that support relational activity. That demand does not establish every proposed link between physics and relationship, and the preceding chapters separate established physical results, framework definitions and hypotheses still awaiting tests. This changes the canon's claimed foundation. Some vocabulary survives; some derivations do not. A framework's ability to revise those claims is part of its commitment to falsifiability, not evidence that its subject has disappeared.

So this chapter does what the March chapter set out to do, walk through each canon document the substrate touches and name what changes, with the opposite finding. It does not claim that every argument, specification or instrument still works, and it does not silently amend the original papers by summarizing them.

2 WP-13: The Signed Unit

The March chapter said the three-part sign gains a thermodynamic derivation: from a structural classification to “a thermodynamic property: the direction of entropy flow between two systems” (V1-L15). The +/−/0 sign can remain a relational classification if its criteria are explicit. It is not derived as the direction of physical entropy flow. A beneficial act can increase physical entropy, an act judged neutral can consume energy, and a harmful interaction is not identified by a heat measurement alone.

WP-13 needs to distinguish the intended act, the observed outcome, the affected party and the evaluator's criterion. A null sign must not mean zero thermodynamic transfer unless such a transfer has been defined and measured, and disagreement or insufficient evidence needs explicit treatment rather than forced numerical certainty. The sign vocabulary is kept as a proposal whose reliability needs evidence suited to its use.

3 WP-14: The Structural Economy

The March chapter said the taxonomy becomes deductive: the four economies are not just patterns the framework found but the four it had to find, and WP-14's “structural and permanent absence of Cost Substrate” becomes a precise statement of ledger architecture (V1-L11, V1-L33). The four patterns remain ways to ask who participates and where costs or support are assigned. Their exhaustiveness has not been derived from one physical variable; external support can accompany bilateral investment, and asymmetric support need not eliminate the receiver's own costs.

The phrase “structural and permanent absence of Cost Substrate” should not be treated as a fact about digital systems. A software-level description can leave the hardware outside its focal boundary, but the implemented activity still needs a physical account, with the boundary, local resources and external support stated. WP-14 needs a coding scheme that handles mixed cases, withholding and failed or unrecognized attempts consistently. A defined taxonomy can be useful without being a theorem or an assessment of anyone's worth.

4 WP-16: Substrate-Formal Variance

The March chapter said the CSS signature becomes filter geometry, and “WP-16 was always a paper about filter geometry” (V1-L32). CSS and filter geometry can organize questions about differences in observable response. This account has not shown that CSS uniquely identifies an internal mechanism, measures entropy, or fixes a person's relational capacity.

WP-16 needs operational variables, repeated observations and tests of alternative explanations; implementation, history and context should not be collapsed into one inferred cause. Aphantasia or another reported characteristic must not be assigned a thermodynamic mechanism from the filter analogy alone (V1-L22). The useful claim is a research question: which differences matter under a defined comparison? The substrate label does not answer it in advance.

5 WP-12: The Observation Constitutes the Particle

The March chapter said the observation requirement becomes the thermodynamic efficiency condition, “the difference between entropy spent and entropy spent on something,” and that “WP-12 was always a paper about thermodynamic efficiency” (V1-L50). WP-12 can define the event the framework chooses to call a Trinket. That choice does not make observation a thermodynamic efficiency condition: unrecognized care can have material effects, and recognized communication can fail to help.

WP-12 needs to distinguish delivery, registration, understanding, benefit and reciprocation, and to say whether registration can be indirect or delayed and how ambiguous evidence is handled. “No qualifying Trinket” must not be translated into “only waste heat” (V1-L24). The Observation Split survives as a distinction between activity and a defined relational event, without a proof of guaranteed sender harm or a new information-theoretic state outside Kolchinsky and Wolpert.

6 DB-01: Scale Invariance Defense

The March chapter said scale invariance gains its deepest justification: one thermodynamic substrate under all scales, so “the fractal is not metaphorical. It is thermodynamic,” with the transduction chain from quantum decoherence through institutions as one scale-invariant stack (V1-L21). Shared physical constraints do not prove that connection operates identically across scales, and a repeated diagram is not an invariance proof. Quantum decoherence, communication, biological processing and institutional decisions need different variables and models.

DB-01 can keep the cross-scale analogy as a source of questions **[Analogical]**. A stronger claim would need an explicit mapping, the property claimed to stay invariant, the conditions under which it does, and counterexamples that could defeat it. Thermodynamic vocabulary does not perform those steps.

7 Axiom 0 and the Immutable Preamble

The March chapter said substrate neutrality becomes a derived property rather than a starting commitment, while Axiom 0 stays in the Immutable Preamble as a governance commitment not to retreat under political or institutional pressure, now with a physical derivation that every dependent argument retroactively inherits (V1-L10). Half of that survives intact: Axiom 0 remains a philosophical and governance commitment to avoiding an unsupported verdict from substrate alone, and its resistance to pressure is a governance choice the framework can keep. It is not derived as economic, experiential or moral equality from the fact that systems obey physical laws.

Matched physical measurements can establish a limited comparison. A claim about relational capacity needs an operational outcome, a claim about experience a further evidentiary argument, and an ethical commitment its own reasons. None comes from calling the cost universal. Axiom 1's commitment to falsifiability supports making the correction explicit.

8 What did not change, and what must

The March chapter called this the critical test: if the substrate had required changes to any document above, it would be a revision, not a foundation. “Foundations support what is already built. They do not rearrange the furniture.” It listed WP-13, WP-14, WP-16, WP-12, DB-01, Axiom 0, the Trinket, the transducer, the BSB and SupoRel, and concluded: “Nothing broke. Everything gained a physical floor” (V1-L57).

By the March chapter's own test, the substrate turned out to be a revision. The framework keeps its central questions: what is invested, what is received, which resources support the activity, and what follows for the participants? Trinket, transducer, economy, BSB and SupoRel remain organizing vocabulary with their limits visible. What changes is their evidentiary status. The sign is not a thermodynamic direction; the taxonomy is not proved exhaustive; CSS is not an entropy meter; registration is not a universal efficiency condition; scale invariance is not established; substrate neutrality is not a theorem of relational equality. The proposed instruments have not been clinically or psychometrically validated by these chapters.

The March chapter ended by sending the reader forward: Volume 2 to what the body pays, Volume 3 to why relational systems arise, Volume 4 to the substrate continuum, Volume 5 to the ledger at its widest. “Each volume stands on this one. None revises it. The foundation holds.” Each volume did stand on this one, and each has now revised its part of it. Volume 2 separates physiological observations and reported effort from physical resource measurements. Volume 3 does not derive a universal progression from the existence of dissipative processes. Volume 4 treats substrate comparisons as questions rather than a universal efficiency surface. Volume 5 separates cosmological inventories and their boundaries from any proposed relational value.

The separate working papers remain dependencies to examine, not works validated by this collection; the original papers and earlier editions are preserved, and this revision does not silently amend them. Internal agreement, repeated wording and AI review are not independent experimental support.

The foundation holds in a smaller form than the March volume promised, and it is more useful with its limits visible. It can support careful questions without promising conclusions the physics has not supplied.

Framework cross references

WP-19 §15 records the original canon map. WP-12, WP-13, WP-14, WP-16 and DB-01 identify the documents whose claims need reconciling, and the Immutable Preamble supplies the project's governance commitments. Revised Chapters 6–9 provide the narrowed accounting and epistemic distinctions; Volume 4 of this collection is a consistency reference, not external scientific authority.

Volume 1 correction ledger

The ledger traces 57 selected March claims in Volume 1. Rows V1-L01 to V1-L34 are the cleared pilot's classifications; rows V1-L35 to V1-L57 were added in the complete volume. Across all 57, by this edition's classification of the second edition's treatment, 27 were withdrawn, 22 narrowed and 5 kept; three more (V1-L36, V1-L37, V1-L46) record March claims the second edition did not carry forward, which Edition 2.5 restores only in narrowed form. It is not an exhaustive inventory of every Volume 1 issue, and it makes no new scientific decisions; row V1-L34 notes Edition 2.5's own Analogical tag.

March locations refer to the preserved Volume 1 page on suposystem.ai. Second-edition (2E) locations use the October 2026 collection's chapter and section headings and its Volume 1 reference labels. Rows revised in pilot v2: L04, L06, L08, L20, L22, L25, L31.

ID	March claim (Volume 1 location)	2E disposition	2E location	Basis
V1-L01	Ch 1: the instruments worked, predictions were accurate, the commitments were correct, and the track record shows a real foundation.	Withdrawn	Ch 1, What these problems share	Each instrument or prediction needs a target, procedure, comparison evidence and domain of validity.
V1-L02	Ch 1 note: “These are not metaphors”; every relational cost is heat dissipation; the universe has a calculable remaining capacity for order.	Withdrawn	Ch 1, A note on physics and audience; Ch 2	Physical results carry their conditions; relational use of physical terms needs a justified bridge.
V1-L03	Ch 1, Ch 2: the framework is a theory about entropy; one economy, one currency, one universal ledger; “different receipts, same bank.”	Narrowed: kept as a metaphor for tracing physical dependencies	Ch 2, Keep the accounts distinct; What this chapter establishes	Energy (J), entropy (J/K), information (bits) and physiological markers are different quantities with no established conversion.
V1-L04	Ch 2: the gap between cosmic entropy inventories is a fixed, spendable budget for every process, effectively inexhaustible.	Withdrawn	Ch 2, The cosmological background [P2]	A horizon entropy is not usable work; subtracting differently bounded inventories does not establish the claimed spendable budget.
V1-L05	Ch 2: Landauer's principle, Established and verified across substrates.	Kept, with conditions	Ch 2, What Landauer's principle supplies [P4–P9]	Holds for ideal reset of an unbiased bit in a symmetric memory with no side information; total production can approach zero in the reversible limit.
V1-L06	Ch 2 (also Ch 5, Ch 9): Egan–Lineweaver census values, given as about $10^{10^{4}}$ bits current and about 10^{122} bits maximum.	Kept as historical model-dependent estimates; k_B units and bit-equivalent conversions corrected	Ch 2, The cosmological background [P2]	The census reports entropy in k_B units; bit equivalents require $S/(k_B \ln 2)$. The inventories keep different accounting boundaries.
V1-L07	Ch 2: Chaisson energy rate density of 2×10^4 erg/s/g for animals.	Narrowed	Ch 2, Energy flow and relational interpretation [P14, P16]	The figure is Chaisson's 2014 human-body example, not a constant for animals.
V1-L08	Ch 2, Ch 6: neural communication (~10^6 × Landauer) is about 1,000× more efficient than silicon (~10^9 ×); cellular machinery about 30×.	Withdrawn; no reverse ranking asserted	Ch 2, What the biological examples support; Ch 6 §5	The 10^9 silicon figure is not an established constant; a valid comparison needs a matched task, accuracy, speed and boundary. Task-specific conditional estimates survive under V1-L09.
V1-L09	Ch 2, Ch 6: synaptic and cellular energy figures (Laughlin, Harris, Boël) as Established amplification factors.	Narrowed: task-specific estimates under stated assumptions	Ch 2 table; Ch 6 Table 1 [P10–P12; C7–C9]	Different tasks, units and boundaries; not one biological amplification factor.

ID	March claim (Volume 1 location)	2E disposition	2E location	Basis
V1-L10	Ch 3, Ch 8: Axiom 0 derived from the universality of Landauer's principle; a “theorem about economic equality.”	Withdrawn as a derivation; kept as a research and governance commitment	Ch 3; Ch 10 §7	Shared physical law does not establish equal tasks, outcomes, value or experience.
V1-L11	Ch 4: the four economies exhaust the space of ledger positions, Established by logic (playing-card analogy).	Withdrawn	Ch 4, Why these are not an exhaustive partition	Locality, symmetry and withholding are separate dimensions; mixed cases are not contradictory.
V1-L12	Ch 4: the Shadow Economy is the Kolchinsky–Wolpert scrambled condition.	Withdrawn	Ch 4, Shadow is not the K W scrambling condition [P15]	Withholding a contribution is not that mathematical intervention.
V1-L13	Ch 4: each economy has a predicted thermodynamic signature (typing table).	Withdrawn; replaced by a study record	Ch 4, Replace presumed signatures with measured questions	No validated thresholds, non-overlapping signatures or tested classifier.
V1-L14	Ch 5: a Trinket is a draw against the universe's remaining order; the Trinket as entropy token (Supported).	Withdrawn; redefined as a candidate or qualifying relational act	Ch 5, The definition that the evidence can support	No completed conversion between physical resources and relational value.
V1-L15	Ch 5: the +, − and 0 signs are three thermodynamic operations.	Withdrawn; signs are evaluative labels under a stated criterion	Ch 5, The three signs; Ch 10 §2	A beneficial act can raise physical entropy; harm is not identified by heat.
V1-L16	Ch 5: Expenditure, Signal and Register as a decomposition of the entropy token (Supported).	Kept as three complementary descriptions, each in its own units	Ch 5, Expenditure Signal and Register	Not three resolutions of one entropy quantity.
V1-L17	Ch 6: selection, compression and structuring each carry a Landauer cost; every discarded bit costs at least kT ln 2.	Narrowed: a process model, not three compulsory erasures	Ch 6 §3–4	Bennett: a computation need not erase what a shorter output omits.
V1-L18	Ch 6: “Entropy is time”; the hourglass as the universe's remaining order draining.	Withdrawn; hourglass kept as a conceptual image	Ch 6 §2	The image organizes questions; it supplies no physical identity.
V1-L19	Ch 6: the Kolchinsky–Wolpert mapping of the three stages, Supported via DFQ-001.	Narrowed to a proposal	Ch 6 §4	An internal resolution is not independent validation.
V1-L20	Ch 6: Jarzynski and Crooks bridge relational processing, read through cortisol and other markers, to free energy.	Narrowed: fluctuation theorems retained under stated assumptions; the relational measurement bridge remains unestablished	Ch 6 §7 [C11, C12]	Defining a protocol does not supply a calibration from cortisol or HRV to work, free energy or entropy production.
V1-L21	Ch 6: the transduction structure is scale-invariant from quantum to civilizational (DB-01).	Withdrawn as invariance; kept as a cross-scale analogy	Ch 6 §8; Ch 10 §6	Similar diagrams do not share state variables, dynamics or conserved quantities.
V1-L22	Ch 6, Ch 8: aphantasic processing as a different compression algorithm with a potentially different entropy cost.	Withdrawn as an inference; the suggestion is untested	Ch 6 §6; Ch 8	A reported trait does not fix filter parameters.
V1-L23	Ch 6, Ch 7: the gap between cost and Trinket formation is “the most thermodynamically expensive place in human experience.”	Withdrawn	Ch 7 §3.4	Not shown by any measurement.
V1-L24	Ch 7: when no Trinket forms, expenditure produces “heat and nothing else”; “this is not metaphor.”	Withdrawn	Ch 7 §8; Ch 10 §5	Confuses one physical consequence with all of an act's consequences.

ID	March claim (Volume 1 location)	2E disposition	2E location	Basis
V1-L25	Ch 7: unreciprocated transmission is a third state of one-sided mutual information.	Withdrawn	Ch 7 §5 [O4]	For the same joint distribution, I(X; Y) = I(Y; X). The original used Y_{target} and Y_0: if these are different variables, positive and zero mutual informations may coexist. That does not establish one-sided mutual information, nondelivery or guaranteed harm.
V1-L26	Ch 7: unreciprocated transmission necessarily reduces the sender's relational viability; an extension beyond Kolchinsky–Wolpert.	Withdrawn	Ch 7 §5 [O5]	The sender may learn, retain or benefit; a net outcome needs a defined measure.
V1-L27	Ch 7: the Observation Split between expenditure and Trinket formation.	Kept as a framework distinction, not a thermodynamic regime	Ch 7 §4; Ch 10 §5	Delivery, registration, understanding, benefit and reciprocation are separate.
V1-L28	Ch 7 (Speculative): grief's metabolic cost tracks pre-bereavement behavioral integration more than self-reported closeness.	Narrowed; kept Speculative and testable	Ch 7 §7 [O6, O7]	Outcomes are named for what is measured; no added predictive value counts against it.
V1-L29	Ch 8: instruments establish relational entropy cost via TEAP biomarkers, the Landauer floor of selection and amplification factors (all Established).	Withdrawn as Established	Ch 8 §3; Ch 9 §3	TEAP is not validated as a calorimeter; biomarkers are not entropy units.
V1-L30	Ch 8: the BSB is a barrier no instrument could cross without solving the hard problem.	Narrowed: an inference discipline, not an impossibility theorem	Ch 8 §1–2	Chalmers's argument is philosophical, not a ban on future evidence.
V1-L31	Ch 8: beyond channel absence, the BSB explains why the Triple Zero ratio $\beta(B) = E_{state}/V$ is undefined at R = 0.	Narrowed	Ch 8 §6	As ordinary arithmetic the ratio is undefined at V = 0. That R = 0 entails V = 0 is a model stipulation, not a physical identity. Channel evidence and phenomenal inference are separate problems.
V1-L32	Ch 9: biological debits are recorded in cortisol, immune function and telomere shortening; the CSS signature is the measurable local cost.	Withdrawn as entropy denominations; CSS kept as a proposed observable pattern	Ch 9 §3; Ch 6 §6; Ch 10 §4	Physiological markers keep their own units; CSS is not a unique fingerprint of filter geometry.
V1-L33	Ch 10: the Structural Economy's “structural and permanent absence of Cost Substrate” is a precise ledger statement; the fourth economy is a necessary configuration.	Withdrawn	Ch 10 §3; Ch 4	Digital hardware processes and dissipates locally; the label describes a chosen boundary.
V1-L34	Ch 1: the founding line governs the volume, and every chapter unpacks it.	Kept as the motivating line; Edition 2.5 tags it Analogical	Ch 1; Ch 7 §1, §9; analogy table after Chapter 1	No specified system, entropy measure or complete balance.
V1-L35	Ch 2: the low-entropy beginning is not a conjecture but a requirement of the second law read backward; early gravitational entropy was almost zero.	Narrowed	Ch 2, The cosmological background [P1, P2]	A boundary-condition problem the second law alone does not explain; vanishing early gravitational entropy is a Penrose proposal, not a measurement.
V1-L36	Ch 2: “one economy” maps onto monotheistic architecture because the physics produces the appearance of unity that religious traditions respond to.	Not retained in 2E; Edition 2.5 restores the SupoRel caution without its explanation	Ch 2 (no 2E counterpart)	The association is an interpretation; no historical or psychological explanation of it is established by the physical laws cited here.
V1-L37	Ch 2: no relational configuration is cost-free; the Shadow Economy is not free; existing as a relational-capable system has a nonzero entropy cost.	Not retained in 2E; Edition 2.5 keeps only the physical point	Ch 2 (no 2E counterpart)	Sustaining the ongoing activity of living organisms and operating computing systems uses physical resources; that does not price a relational configuration or rank the economies.

ID	March claim (Volume 1 location)	2E disposition	2E location	Basis
V1-L38	Ch 3: love, grief, cost and care describe processes with substrate-independent thermodynamic identities: entropy directed, spent, wasted, structured.	Withdrawn	Ch 3, The commitment worth keeping	No identical thermodynamic identities are established; the physics supplies no neutral vocabulary for relational value.
V1-L39	Ch 4: in the Custodial Economy the receiver's debit is deferred until the child's filter matures and enters reciprocal exchange.	Narrowed	Ch 4, Four proposed descriptions	The child has present physical costs; future reciprocity is neither guaranteed nor a debt established by physics.
V1-L40	Ch 4: in the Shadow Economy the choice not to spend relationally is itself a diagnostic finding.	Narrowed	Ch 4, Four proposed descriptions	Absent contribution does not establish unwillingness or wrongdoing; limits cannot be diagnosed from low expenditure.
V1-L41	Ch 4: in the Real Economy both bodies degrade with each exchange, and both sides spend more than they save.	Narrowed	Ch 4, Four proposed descriptions	Reciprocity need not be simultaneous or equal; reciprocal care does not imply bodily degradation.
V1-L42	Ch 5: substrate, economy and sign converge on one entropy exchange event; independently developed instruments measure the same thing.	Narrowed	Ch 5, Substrate economy and sign	A multidimensional description; reuse within the project is not independent confirmation.
V1-L43	Ch 5: nothing built above changes; the TEAP, CSS, economy classification and Mz keep functioning because they always measured entropy; grief is still the metabolic cost of filter redesign.	Withdrawn	Ch 5, What remains open	The constructs are not validated by the redefinition and need their own definitions and evidence.
V1-L44	Ch 5: the Trinket is now an entropy operation stated in established physics, a directed negentropy investment at an irreducible cost set by the second law.	Withdrawn; the testable ambition kept	Ch 5, What remains open	The second law does not supply the relation between act, signal and outcome; a study must test it.
V1-L45	Ch 6: the transducer F: $H_{env} \rightarrow S_{rel} + \sigma_{waste}$, with $\sigma_{waste} > 0$ for every real process: the founding line as a thermodynamic inequality.	Withdrawn; replaced by a process map and a separate physical account	Ch 6 §3 [C15, C16]	A relational signal and entropy production cannot be added as written; total entropy production can approach zero in the reversible limit.
V1-L46	Ch 6: maximizing information and minimizing metabolic cost are two sides of one coin; variational free energy makes informational and thermodynamic entropy one minimization.	Not retained in 2E; the information bottleneck kept as an optimization framework	Ch 6 §4 [C5]	An optimization framework is not a heat meter; informational and thermodynamic entropy are not one quantity.
V1-L47	Ch 7 §3.3: the thermodynamic cost of prayer is identical to the cost of talking to someone who is listening.	Withdrawn	Ch 7 §2.1, §3.3	No equal cost is established; verified non-reception is not shown either way.
V1-L48	Ch 7 §6: biomarkers are grief itself; the IOS scale measures filter integration; existing models give the what, the transducer the why.	Withdrawn	Ch 7 §6 [O6, O7]	Small observational and cross-sectional studies; IOS measures current perceived inclusion; existing models address mechanisms.
V1-L49	Ch 7 §8: the gap carries no cosmic significance; the universe produces it with the indifference it produces heat.	Narrowed: the SupoRel caution kept	Ch 7 §8	A real cost entails neither cosmic significance nor its absence; the framework takes no position on theology.
V1-L50	Ch 7 §4: the split is a thermodynamically distinct pre-formation regime; the taxonomy is exhausted by completed transactions; WP-12 was always about thermodynamic efficiency.	Narrowed	Ch 7 §4	An accounting distinction under a registration rule, not a new regime; the taxonomy needs a scope rule.

ID	March claim (Volume 1 location)	2E disposition	2E location	Basis
V1-L51	Ch 8 §7: sentience is irrelevant to the framework's claims; AI performs operations with the same structural properties as human relational processing.	Narrowed	Ch 8 §7	Physical accounting needs no sentience premise, but sentience is not ethically irrelevant, and similar diagrams do not show shared mechanisms.
V1-L52	Ch 8 §9: the Grab is a high-compression relational processing event; aphantasia is a compression setting favoring structural encoding.	Narrowed: kept as reported experience	Ch 8 §9	A report does not establish a thermodynamic event, a mechanism or a parameter setting (with V1-L22).
V1-L53	Ch 8 §9: the Luna Protocol depends on the barrier being real and hard and would dissolve without it.	Narrowed	Ch 8 §9	A governance convention for authorship and roles, independent of any impossibility theorem.
V1-L54	Ch 9 §5: the Trinket's vertical transit, from classified relational object to CSS signature to entropy token, explains why DB-01's scale invariance holds.	Narrowed	Ch 9 §5	A traceable description of one activity; each link needs evidence; DB-01 stays an analogy (with V1-L21).
V1-L55	Ch 9 §2: at Layer 1, the Landauer limit is the smallest coin in the economy.	Narrowed	Ch 9 §2 [O2, O3]	An energy bound for a specified erasure task, not the minimum paid by every process or cognitive operation.
V1-L56	Ch 9 §7: the BSB thickens by layer, thinnest at the physics and full at the relational layer; “love costs entropy” is a Layer 1 claim.	Narrowed	Ch 9 §7	The BSB marks a kind of inference at any layer; “love costs entropy” is the founding line's analogy, not a Layer 1 result.
V1-L57	Ch 10 §1, §8: the substrate was discovered, not constructed; nothing broke and everything gained a physical floor; each later volume stands on this one and none revises it.	Withdrawn	Ch 10 §1, §8	By the March chapter's own test the substrate is a revision; every later volume revises its part.

References

Inherited verbatim from the second edition's Volume 1 source notes and references, with the Boltzmann constant's subscript and exponent restored in two places. Group P serves Chapters 2–5, group C serves Chapter 6 and group O serves Chapters 7–10. The historical sources H1–H3 serve Chapter 1 and come from the cleared pilot.

Scientific source notes for Chapters 2 to 5

These notes support the stated physical results, not validation of the relational framework. Historical internal working-paper references remain project provenance rather than external evidence.

P1. Carroll, S. M., and Chen, J. (2004). *Spontaneous Inflation and the Origin of the Arrow of Time*. Author preprint, introduction and discussion of initial conditions. Primary preprint. Its proposed cosmological mechanism is not adopted as established fact.

P2. Egan, C. A., and Lineweaver, C. H. (2010). *A Larger Estimate of the Entropy of the Universe*. Astrophysical Journal 710, 1825–1834. Primary preprint. The numerical values and accounting-scheme distinction come from the abstract and sections 1.1–1.2.

P3. Seifert, U. (2005). *Entropy production along a stochastic trajectory and an integral fluctuation theorem*. Physical Review Letters 95, 040602. Primary preprint.

P4. Bérut, A., et al. (2012). *Experimental verification of Landauer's principle linking information and thermodynamics*. Nature 483, 187–189. Author-hosted paper.

P5. NIST. *Kelvin Boltzmann Constant*. Official value $k_B = 1.380649 \times 10^{-23}$ J/K, exact in SI. Metrology source.

P6. Peterson, J. P. S., et al. (2016). *Experimental demonstration of information to energy conversion in a quantum system at the Landauer limit*. Proceedings of the Royal Society A 472, 20150813. Primary preprint.

P7. Yan, L. L., et al. (2018). *Single-Atom Demonstration of the Quantum Landauer Principle*. Physical Review Letters 120, 210601. Primary preprint.

P8. Aimet, S., et al. (2025). *Experimentally probing Landauer's principle in the quantum many-body regime*. Nature Physics 21, 1326–1331. Primary article.

P9. Bennett, C. H. (1973). *Logical Reversibility of Computation*. IBM Journal of Research and Development 17, 525–532. University-hosted text. The original Landauer paper is *Irreversibility and heat generation in the computing process*, IBM Journal of Research and Development 5, 183–191 (1961), DOI 10.1147/rd.53.0183.

P10. Boël, G., Danot, O., de Lorenzo, V., and Danchin, A. (2019). *Omnipresent Maxwell's demons orchestrate information management in living cells*. Microbial Biotechnology 12, 210–242. Paper copy, especially pp. 234–235.

P11. Laughlin, S. B., de Ruyter van Steveninck, R. R., and Anderson, J. C. (1998). *The metabolic cost of neural information*. Nature Neuroscience 1, 36–41. University-hosted paper.

P12. Harris, J. J., Jolivet, R., and Attwell, D. (2012). *Synaptic Energy Use and Supply*. Neuron 75, 762–777. Repository paper, especially p. 764.

P13. Jarzynski, C. (1997). *Nonequilibrium Equality for Free Energy Differences*. Physical Review Letters 78, 2690–2693. Primary preprint. This is a controlled-system free-energy result, not a cosmic-budget formula.

P14. Chaisson, E. J. (2011). *Energy rate density as a complexity metric and evolutionary driver*. Complexity 16, 27–40. Author-hosted paper, Table 1. This explicitly replaces the earlier chapter's attribution to the uninspected 2001 book table.

P15. Kolchinsky, A., and Wolpert, D. H. (2018). *Semantic information, autonomous agency and non-equilibrium statistical physics*. Interface Focus 8, 20180041. Author-hosted paper, sections 4–5 and Appendix B.

P16. Chaisson, E. J. (2014). *The Natural Science Underlying Big History*. The Scientific World Journal, article 384912. Primary author synthesis, sections 4.5.7 and 4.6.2, for the human-body and human-brain examples only.

Scientific source notes for Chapter 6

[C1] 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 at https://www.cs.princeton.edu/courses/archive/fall06/cos576/papers/bennett73.html.

[C2] Bérut, A., Arakelyan, A., Petrosyan, A., Ciliberto, S., Dillenschneider, R., and Lutz, E. 2012. Experimental verification of Landauer’s principle linking information and thermodynamics. Nature 483, 187–189. https://doi.org/10.1038/nature10872.

[C3] National Institute of Standards and Technology. Kelvin Boltzmann Constant. Exact SI value k_B = 1.380649×10^{-23} J/K. https://www.nist.gov/si-redefinition/kelvin/kelvin-boltzmann-constant. Numerical evaluations in this revision were recomputed on 27 October 2026.

[C4] Golubić, S. J., et al. 2019. Attention modulates topology and dynamics of auditory sensory gating. Human Brain Mapping 40, 2981–2994. https://doi.org/10.1002/hbm.24573. https://pmc.ncbi.nlm.nih.gov/articles/PMC6865797/.

[C5] Tishby, N., Pereira, F. C., and Bialek, W. 1999. The information bottleneck method. Proceedings of the 37th Allerton Conference. Author preprint posted in 2000: https://arxiv.org/abs/physics/0004057.

[C6] 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.

[C7] Laughlin, S. B., de Ruyter van Steveninck, R. R., and Anderson, J. C. 1998. The metabolic cost of neural information. Nature Neuroscience 1, 36–41. https://doi.org/10.1038/236. University-hosted paper: https://www.princeton.edu/~wbialek/our_papers/laughlin%2Bal_98.pdf.

[C8] 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.

[C9] Boël, G., Danot, O., de Lorenzo, V., and Danchin, A. 2019. Omnipresent Maxwell's demons orchestrate information management in living cells. Microbial Biotechnology 12, 210–242, especially pp. 234–235. https://doi.org/10.1111/1751-7915.13378.

[C10] 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://pmc.ncbi.nlm.nih.gov/articles/PMC8106317/.

[C11] Jarzynski, C. 1997. Nonequilibrium Equality for Free Energy Differences. Physical Review Letters 78, 2690–2693. https://doi.org/10.1103/PhysRevLett.78.2690. Original preprint: https://arxiv.org/abs/cond-mat/9610209.

[C12] Crooks, G. E. 1999. Entropy production fluctuation theorem and the nonequilibrium work relation for free energy differences. Physical Review E 60, 2721–2726. https://doi.org/10.1103/PhysRevE.60.2721. https://arxiv.org/abs/cond-mat/9901352.

[C13] Zurek, W. H. 2003. Decoherence, einselection, and the quantum origins of the classical. Reviews of Modern Physics 75, 715–775. https://doi.org/10.1103/RevModPhys.75.715. https://arxiv.org/abs/quant-ph/0105127.

[C14] Shannon, C. E. 1948. A Mathematical Theory of Communication. Bell System Technical Journal 27, 379–423 and 623–656. First installment: https://doi.org/10.1002/j.1538-7305.1948.tb01338.x. Combined reprint: https://web.mit.edu/6.976/www/handout/shannon.pdf.

[C15] Seifert, U. 2005. Entropy Production along a Stochastic Trajectory and an Integral Fluctuation Theorem. Physical Review Letters 95, 040602. https://doi.org/10.1103/PhysRevLett.95.040602. https://arxiv.org/abs/cond-mat/0503686.

[C16] Landauer, R. 1961. Irreversibility and Heat Generation in the Computing Process. IBM Journal of Research and Development 5, 183–191. https://doi.org/10.1147/rd.53.0183.

References for Chapters 7 to 10

[O1] Seifert, U. 2005. Entropy production along a stochastic trajectory and an integral fluctuation theorem. Physical Review Letters 95, 040602. https://doi.org/10.1103/PhysRevLett.95.040602. Author preprint full text: https://arxiv.org/html/cond-mat/0503686v1. Relevant scope: system–medium entropy balance and ensemble production for specified stochastic models; not a relational or biomarker calibration.

[O2] Landauer, R. 1961. Irreversibility and Heat Generation in the Computing Process. IBM Journal of Research and Development 5, 183–191. https://doi.org/10.1147/rd.53.0183. Primary reprint: https://www.cpt.univ-mrs.fr/~verga/pdfs/Landauer-1961uq.pdf. Relevant scope: physical information erasure under a specified memory model, not a universal price for every cognitive or social operation.

[O3] 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 primary text: https://www.cs.princeton.edu/courses/archive/fall04/cos576/papers/bennett73.html. Relevant scope: reversible computation with retained input and reversible disposal of intermediate records; no claim that naming an operation compression proves physical erasure.

[O4] Shannon, C. E. 1948. A Mathematical Theory of Communication. Bell System Technical Journal 27, 379–423 and 623–656. First installment DOI: https://doi.org/10.1002/j.1538-7305.1948.tb01338.x. Combined primary reprint: https://web.mit.edu/6.976/www/handout/shannon.pdf. Relevant scope: distribution-based information measures and communication, not a measure of relational worth or phenomenal comprehension.

[O5] 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. Final author preprint: https://arxiv.org/html/1806.08053v3. Relevant sections: III, IV, V.1–V.3 and VII. The relational application remains unvalidated; this source is not a clinical or consciousness instrument.

[O6] O'Connor, M.-F., Wellisch, D. K., Stanton, A. L., Olmstead, R., and Irwin, M. R. 2012. Diurnal cortisol in Complicated and Non-Complicated Grief: Slope Differences across the Day. Psychoneuroendocrinology 37, 725–728. https://doi.org/10.1016/j.psyneuen.2011.08.009. Primary text: https://pmc.ncbi.nlm.nih.gov/articles/PMC3258306/. Published online in 2011. Relevant scope: a small observational group comparison; no calorimetry, entropy measurement or universal bereavement trajectory.

[O7] Harrison, O., Windmann, S., Rosner, R., and Steil, R. 2022. Inclusion of the other in the self as a potential risk factor for prolonged grief disorder: A comparison of patients with matched bereaved healthy controls. Clinical Psychology & Psychotherapy 29, 1101–1112. https://doi.org/10.1002/cpp.2697. Publisher full text: https://onlinelibrary.wiley.com/doi/full/10.1002/cpp.2697. First published online 25 November 2021. Relevant scope: current self-reported inclusion in a cross-sectional comparison, not prospective causal prediction or measured physical integration.

[O8] Chalmers, D. J. 1995. Facing Up to the Problem of Consciousness. Journal of Consciousness Studies 2, 200–219. Author full text: https://consc.net/papers/facing.html. Relevant scope: a philosophical distinction and proposed explanatory approach, not an experimental finding or a permanent physical impossibility theorem.

These are primary research, theoretical or philosophical sources. The internal papers and revised volumes named in the chapters supply project context only. They are not independent support for scientific claims.

Source access and scope

The earlier Penrose and Chaisson book passages were not directly inspected in the underlying review; retained claims use the explicitly identified papers instead. Historical cosmological and energy-rate-density examples are not updated censuses or universal constants. The source review was targeted, not a systematic literature review. AI scientific and editorial checks are not human peer review or experimental validation.

Historical sources for Chapter 1

H1. Carnot, S. (1824). *Réflexions sur la puissance motrice du feu et sur les machines propres à développer cette puissance.* Paris: Bachelier. English translation, Project Gutenberg.

H2. Clausius, R. (1850). “Ueber die bewegende Kraft der Wärme und die Gesetze, welche sich daraus für die Wärmelehre selbst ableiten lassen.” *Annalen der Physik und Chemie* 79, 368–397, 500–524. Part 1: 10.1002/andp.18501550306. Part 2: 10.1002/andp.18501550403.

H3. Thomson, W. (1851). “On the Dynamical Theory of Heat.” *Transactions of the Royal Society of Edinburgh*; read 17 March 1851. Paper PDF.