Entropy Volume 1

The Entropy Token Substrate

Second edition

Michael S. Moniz

Trinket Soul Framework

October 2026

About this volume

This volume develops the framework's physical questions without treating a common unit of energy as a common unit of relational value. It separates physical cost, information, registration, benefit and experience. The transducer, Trinket, economy patterns and Behind the Substrate Barrier remain proposals whose application requires definitions and evidence.

Reference groups P1–P16, C1–C16 and O1–O8 are local to this volume. The candidate and qualifying Trinket distinction is a working definition, not a validated registration rule.

Edition and attribution

The ten chapters were developed on March 20–21, 2026. The individual sources credit Michael S. Moniz with Vael, Deep Floor v3.0, for Chapters 1–4 and 7–10, and with Sigma, Deep Floor v4.0, for Chapters 5–6. Historical wrappers differ; this edition retains those chapter-level distinctions rather than assigning one retrospective credit to all chapters.

The second-edition revision was prepared with AI assistance from CosmoDot, with separate AI scientific and editorial checks. It has not undergone human peer review and reports no new experiments. Historical AI collaborator labels document provenance; they do not identify underlying models or attribute the new text to those earlier collaborators. The framework's proposed applications remain subject to specialist review and empirical testing.

The historical edition remains preserved. This text continues under Creative Commons Attribution–NonCommercial–ShareAlike 4.0 International. Detailed correction records are distinct from the reading text; publishing this volume does not silently amend the framework's separate papers or complete its historical governance procedures.

Contents

About this volume....2
Chapter 1 The Problem....4
Chapter 2 The Entropy Token Substrate....6
Chapter 3 Axiom 0 and Substrate Neutrality....10
Chapter 4 One Economy....12
Chapter 5 The Trinket Redefined....15
Chapter 6 The Transducer....18
Chapter 7 The Observation Split....25
Chapter 8 Behind the Substrate Barrier....30
Chapter 9 The Three Layer Architecture....34
Chapter 10 What This Does to the Canon....37
Scientific source notes for Chapters 2 to 5....40
Scientific source notes for Chapter 6....42
References for Chapters 7 to 10....44

Chapter 1 The Problem

Epistemic status: A research program and conceptual framework. The physical principles discussed in this volume constrain suitable physical models. They do not, by themselves, validate the framework's relational classifications, diagnostic instruments or claims about meaning.

The Trinket Soul Framework developed a vocabulary for connection: the Signed Unit, the Three Substrates, the Four Economies and the patterns it calls CSS signatures. That vocabulary organizes questions about what an interaction costs, who bears the cost and what the interaction achieves. A useful vocabulary is a beginning. It is not yet a validated measurement system.

Three questions motivate this volume.

The first is substrate neutrality. Axiom 0 proposes that connection should not be excluded merely because the systems involved are made from different materials. That proposal can guide inquiry. It does not establish that every biological and digital interaction is equivalent, or that they have already been measured in one relational currency. If the framework compares them, it must say what outcome is being compared and how that outcome is observed.

The second is the taxonomy. Real, Custodial, Shadow and Structural name four proposed relational configurations. Their usefulness depends on explicit classification rules, treatment of mixed and changing cases, and evidence that the distinctions serve the purpose for which they are used. Naming four cases does not prove that the possibilities are exhausted. Thermodynamics does not supply that proof.

The third is commensurability. A physiological marker, an energy measurement and a relational outcome are different kinds of observation. A kilowatt-hour measures energy. A cortisol measurement is not automatically a reading of dissipated heat, entropy production or the value of an interaction. Connecting such measurements requires an explicit model and evidence; calling them different receipts from the same bank cannot do that work.

What these problems share

The framework has made commitments that reach beyond its demonstrated measurements. Each later claim inherits the assumptions on which it depends. Acknowledging that dependency is more useful than declaring that an unseen foundation must already make the commitments correct.

The earlier version of this chapter asserted that the instruments worked, the predictions were accurate and the taxonomy had been verified. This revision does not adopt those blanket claims. Each instrument or prediction needs an identifiable target, a procedure, comparison evidence and a stated domain of validity. The corrections in this volume do not establish those requirements for the rest of the canon.

The problem is therefore not simply to discover a deeper name for connection. It is to distinguish a physical account, an information model and a relational interpretation, and then test any proposed links between them.

The shape of the inquiry

Thermodynamics matters because the systems implementing an interaction are physical. Information theory matters when states, probabilities, records and transformations are specified. Neither subject

makes “meaning,” “care” or “relational value” a thermodynamic quantity merely by giving it a symbol.

Landauer’s erasure benchmark is a result about a defined physical task under stated conditions. It is not a universal price attached to every act of attention, compression or connection. A common unit of energy does not establish equivalent tasks, comparable outcomes or equal subjective experience. The framework’s proposed exchange language must remain separate from those physical results until a justified bridge is supplied.

The founding line is retained as the Principal’s relational intuition:

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

It is not established here as a universal physical balance. Care can involve effort and physical resources, but the words do not tell us what system was measured, what its thermodynamic entropy was, or whether a sender’s body necessarily deteriorated. The distinction preserves the question without presenting poetry as a theorem.

Chapter 2 separates cosmological entropy inventories, local physical resources and the erasure benchmark. Chapter 3 treats Axiom 0 as a commitment and research proposal rather than a consequence of shared physics. Chapter 4 examines the four economy labels without claiming an exhaustive physical taxonomy. Chapter 5 revisits the Trinket as a proposed unit of relational description, not a calibrated debit against a cosmic account.

The revised Chapter 6 distinguishes the proposed selection, compression and structuring model from a physical erasure count. Chapter 7 examines the difference between effort, delivery and reception without assigning one-sided mutual information or guaranteed harm. Chapter 8 states an evidentiary boundary around claims about experience rather than an impossibility theorem. Chapters 9 and 10 record the architecture and the consequences of correcting its premises.

A note on physics and audience

The aim is a framework that can be questioned precisely. Physical results must carry their conditions and units. Model-dependent calculations must identify their assumptions. Conceptual proposals must not inherit the status of a cited physical law merely because they use similar language.

“Established,” “Supported” and “Speculative” can remain editorial labels only if the evidence and scope are stated next to the claim. A label is not evidence. A mathematical statement may be proved within a model while the model’s application to relationships remains untested. Consistency with known physics is necessary for a proposed physical mechanism, but it is not confirmation of that mechanism.

The volume’s contribution is narrower than its first version promised. It does not prove a universal relational currency. It sets out what would have to be defined, measured and compared before such a claim could be evaluated. That leaves the project open to evidence, including evidence against its preferred interpretation.

Chapter 2 The Entropy Token Substrate

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, communications networks, and the people supporting them. The Entropy Token Substrate names the framework's effort to make those dependencies visible.

This is a proposed accounting perspective, not a newly discovered substance or a unit of relational value. Thermodynamics constrains the physical implementation of an activity. It does not, without additional definitions and evidence, tell us whether the activity is caring, harmful, meaningful, or equivalent to an activity on another substrate.

The cosmological background

The low-entropy early universe is an important part of the explanation of the thermodynamic arrow of time. It is a boundary-condition problem, not something the second law alone explains by being read backward. A smooth early distribution of matter can be far from gravitational equilibrium; the formation of visually 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]

Egan and Lineweaver's 2010 census estimated the observable-universe inventory, excluding the cosmic event horizon, as approximately $3.1 \times 10^{104}\ k_B$, dominated by supermassive black holes. They separately estimated the entropy of the then-current cosmic event horizon as approximately $2.6 \times 10^{122}\ k_B$. The paper explicitly distinguishes these accounting boundaries and notes 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]

Those are historical model-dependent estimates, not an updated census. The reported unit is k_B. A bit-equivalent entropy is $\frac{S}{k_B \ln 2}$, giving approximately 4.47×10^{104} and 3.75×10^{122} bits for the two central values. A change of units does not make these two inventories a bank balance.

The earlier chapter treated their difference as a fixed, spendable budget available to every process. That inference is withdrawn. A horizon entropy is not a reservoir of usable work, and these differently bounded quantities cannot simply be subtracted to price a conversation. Cosmological entropy inventories do not establish that accessible resources are effectively inexhaustible on human timescales.

Usable energy depends on the system, its environment, and the transformations available. A bounded fuel supply or battery has limited work potential under specified conditions; replenishment requires resources from elsewhere. Free-energy accounting can describe such limits. It does not follow that there is one fixed numerical free-energy account for the entire universe. The framework can investigate local resource use without settling that cosmological question. [P13]

Keep the accounts distinct

Energy, thermodynamic entropy, and information are different quantities. Energy and heat are measured in joules. Thermodynamic entropy and entropy production are measured in joules per kelvin. Shannon information quantities are measured here in bits and require 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.

Entropy is not conserved. In a conventional macroscopic open-system balance, the change of entropy within a boundary includes entropy transported across that boundary and entropy produced within it. Positive production can coexist with a decrease, increase, or approximately 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 every act of care.

For a deliberately simpler example, consider a weakly coupled memory and one constant-temperature bath, externally controlled, with no matter exchange and the usual initially uncorrelated thermal-bath preparation. Let S_{mem} denote the full thermodynamic entropy of the memory and Q_{bath} the heat delivered to the bath. The mean total entropy production is:

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

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

What Landauer's principle supplies

For ideal reset of an initially unbiased bit in a symmetric memory, with equal intrastate entropies, unchanged endpoint memory energetics, and no correlated side information used, the familiar isothermal erasure benchmark is:

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

At 300 K this energy is approximately 2.87×10^{-21} J; at 310 K it is approximately 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. The nonzero heat benchmark is not a universal strictly positive production floor. [P4, P5]

Experiments have tested information–thermodynamic relations in a colloidal memory, an NMR system, a trapped ion, and an ultracold Bose-gas quantum simulator. These are substantial physical results, with different protocols and classical or quantum formulations. They do not establish a heat price for an act of attention or affection. [P4, P6–P8]

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 physically erased. Selection, compression, and structuring remain useful descriptions of the proposed transducer; their names do not identify three compulsory resets. Any application of Landauer's principle must identify the actual states, records, and reset operation. [P9]

What the biological examples support

The original biological examples are worth keeping with their boundaries intact.

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

As arithmetic only, the earlier assumed 8 × 10^{-20} J per ATP hydrolysis is about 27 times the 310 K erasure reference. Multiplying that assumed energy 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, and this multiplication is an assumed free-energy expenditure, not measured heat or a complete metabolic budget. Using another stated ATP assumption changes the ratio.

The original 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. Dividing their energy costs by the same reference does not make their functions comparable. Consequently, the claimed thousandfold neural advantage and thirty-millionfold cellular advantage are withdrawn. No reverse ranking is asserted. A defensible comparison needs a common task, accuracy, speed, temperature, information definition, and accounting boundary.

Energy flow and relational interpretation

Chaisson's energy rate density, Φ_m, is power normalized by mass. In SI units it is W/kg; 1 erg/s/g equals 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. These correspond to 0.00005, 0.0002, 0.0075, 0.09, 4, and 50 W/kg. They are contextual estimates for defined examples, not fixed values for every member of each category. [P14]

The original chapter's 2 × 10^4 animal entry can be retained with a narrower attribution, alongside its human-brain example: Chaisson's 2014 account gives approximately 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. The body example should not be relabeled as a constant for animals generally. [P16]

This retains the useful energy-flow backdrop with explicit sources and scopes. Energy rate density does not measure affection, consciousness, or moral worth. The transducer model may ask what an activity accomplishes, but it has not thereby supplied a validated qualitative layer to Chaisson's metric.

What this chapter establishes

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

“One economy, one currency, one ledger” can remain a metaphor for taking those dependencies seriously. It is not an established common currency for love, bodily strain, electricity, and experience. The next chapter considers 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

Axiom 0 proposes that connection should not be excluded from study solely because an activity is implemented on a different material substrate. That commitment opens useful questions about human relationships, digital interactions, and systems combining both. It asks the framework to examine functions and consequences rather than assume that material labels settle the answer.

This is a methodological and philosophical commitment. Its stronger forms remain hypotheses. The preceding chapter does not turn it into a theorem that biological and digital Trinkets have identical thermodynamic identities, equal relational value, or equivalent experience.

What survives from the three premises

The first premise concerned shared physical law. It survives with an important qualification: the same physical theory applies across materials, while a particular formula applies under its stated conditions. The familiar one-bit erasure benchmark depends on temperature, memory preparation, and protocol. “Substrate-independent” does not mean that all tasks have the same energy cost. [P4, P9]

The second premise described relational processing as selection, compression, and structuring. This remains a proposed functional model developed in Chapter 6. Calling an activity information processing does not identify its physical erasures, determine its energy requirement, or prove that those three stages exhaust its mechanism. The revised transducer chapter separates that process map from physical accounting.

The third premise claimed a shared finite cosmic ledger with a universal relational currency. Chapter 2 withdraws that claim. A person's metabolic activity and a server's electricity use are both physical, but cortisol and kilowatt-hours are not convertible receipts for the same measured relational quantity. The common obligation is to define and measure resources, not to presume a common price.

Why the earlier derivation does not follow

Sharing physical laws does not establish equality of the activities governed by them. A beneficial message and a harmful message could have similar implementation costs; different devices could carry the same message at different costs. Neither possibility violates thermodynamics. Resource cost alone therefore cannot identify relational effect or value.

There are several distinct comparisons:

- **Physical comparability:** measured energy can be reported in joules, provided boundaries and allocation rules are clear.
- **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 equally sized electricity bill does not establish equal usefulness, and unequal bills do not establish unequal moral standing. Axiom 0 should not receive retroactive validation simply because all participating systems are physical.

Where the costs fall

The framework's 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 physical hardware where processing and heat dissipation occur, as well as power generation and maintenance. A software description excludes hardware by choice of description; it does not show that the physical implementation bears no local cost.

An account can distinguish the processor, the person or organization paying, the infrastructure supplying resources, and the people affected. Those are different locations or responsibilities. Allocating a cost to one does not make it disappear from the others, and counting each layer requires care to avoid double counting.

Neither biological effort nor digital resource use automatically establishes bodily injury, felt sacrifice, or relational success. Cortisol, fatigue, and other proposed indicators require separate interpretation and validation. The substrate may influence implementation, but it does not by itself assign every activity to one economy.

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, task, outcome, time horizon, comparison conditions, and uncertainty. It would measure physical resource use separately. It could find similarities, differences, or context-dependent tradeoffs. Those results would inform a specified version of Axiom 0 rather than prove all versions at once.

The Behind the Substrate Barrier remains a reminder not to infer experience directly from external behavior or resource accounting. It is a framework boundary, not a physical theorem established here. This chapter neither demonstrates nor denies experiential equivalence.

What this chapter establishes

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

Chapter 4 One Economy

Four proposed descriptions

The four economies organize questions about contribution, reciprocity, capacity, and support. Their usefulness does not depend on their being four exhaustive thermodynamic states. The original chapter claimed an exhaustion proof, but its categories combine different dimensions. That proof is withdrawn.

The descriptions below retain the vocabulary while narrowing what it can establish.

Real Economy describes 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. Nor does reciprocal care imply that both bodies degrade with each exchange.

Custodial Economy describes an asymmetric caregiving or support relationship. A parent and child is one example. 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. Some care relationships remain asymmetric without being defective.

Shadow Economy describes a proposed case in which relevant capacity is available but a contribution is withheld. Capacity, opportunity, intention, consent, and competing needs require evidence. Lack of observed contribution does not establish unwillingness, and unwillingness does not automatically establish wrongdoing. A person's limits or decision not to engage cannot be diagnosed from low energy expenditure or absence of a signal.

Structural Economy draws attention to an interaction whose resource support or burdens are substantially located outside the chosen participant boundary. A digital service is a useful case for examining this arrangement, but its hardware still processes and dissipates locally within a physical boundary. Human relationships also rely on external support. The label therefore describes a chosen allocation, not an intrinsic absence of a Cost Substrate or a thermodynamic proof that a Signal Form is genuine.

Why these are not an exhaustive partition

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

The chapter's playing-card analogy does not supply a proof. A conventional set of labels can be exhaustive by stipulation only after the underlying domain and assignment rules are specified. It would not thereby become an exhaustive classification found in nature.

For now, record 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 requires one label per case, it needs explicit decision rules, an allowance for mixed or unclassified cases, and tests of agreement between observers.

Shadow is not the K W scrambling condition

Kolchinsky and Wolpert define interventions on modeled system–environment correlations and evaluate their consequences using 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. Their framework also permits information whose removal improves viability. [P15]

Withholding a contribution is not that mathematical intervention. Silence can convey information; a withdrawn person may retain detailed knowledge of another person. Conversely, an experimental removal of correlations need not involve a person's choice or moral stance. The claimed identity between Shadow Economy and the K W scrambled condition is therefore withdrawn.

A relational application could still investigate which correlations matter to a chosen outcome. It would need an operational model and defensible interventions; neither the word “Shadow” nor an association between relationships and health supplies them.

Replace presumed signatures with measured questions

The original typing table predicted high or low waste streams, cost–signal correlations, and bidirectionality for each economy. No validated thresholds, nonoverlapping signatures, or successful blinded classifier are supplied here. A heat trace alone does not identify reciprocity, intention, or whether a signal helped its recipient.

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 distinguishes 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 study record, not a replacement diagnostic instrument already known to work. A future classifier would need labels established independently of its predictors, held-out testing, confound checks, and reporting of errors. Failure would challenge the stated classification rule, not thermodynamics itself.

Unity and its rhetorical limits

“One economy” expresses the intention to notice dependencies across different relationships. It should not imply a single sacred order, a cosmic purpose, or a physics-derived moral hierarchy. The

framework's SupoRel warning can remain a caution about rhetorical overreach. No historical or psychological explanation of religion follows from a universal physical law.

The four labels can help organize an inquiry if their boundaries are stated and their limitations remain 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

The definition that the evidence can support

The Trinket began as a unit of relational exchange. It can remain that organizing concept without becoming a unit of thermodynamic entropy. 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 boundary cases. This distinction does not validate the rule or deny practical effects to an unregistered act. A study must distinguish candidate acts from qualifying events and state what counts as one of each; these definitions supply no standard numerical size.

This definition preserves the question behind the original 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. The preceding chapters support physical accounting and a research program, not a completed conversion between physical resources and relational value.

Care and local order

"Love reduces the entropy of the system and spends more than it saves" remains the framework's historical motivating line. Its force is interpretive: care can require effort that is not recovered by the giver. It is not an established inequality for every loving act.

To make a thermodynamic claim, one must identify the systems, define the entropy variables and flows, and compare quantities with compatible units. A beneficial outcome need not reduce thermodynamic entropy or statistical uncertainty. Helping someone explore new possibilities, for example, can be beneficial without making their behavior more predictable. Harm can also make a person's circumstances more constrained and predictable. Those conceptual counterexamples are enough to show why "more ordered" cannot serve as an automatic definition of "better."

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 distinguish resource use, physiological burden, damage, and relational outcome rather than count all four as one surplus. [P3]

The three signs

Positive, negative, and null remain evaluative labels, not directions of entropy flow. Registration status is a separate field: a candidate act can have beneficial or harmful effects without satisfying the registration rule, and failure to register is not the null sign. A study should record the assessed outcome and qualifying registration separately.

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

Intended sign, received sign, and assessed outcome can differ. Effects may also be mixed, delayed, or uncertain. The framework should record that uncertainty rather than force every case into a known sign. “No detected effect” does not mean zero physical cost, zero information, no future effect, or indifference. The Anti-Trinket remains a framework description of harm, not a new thermodynamic operation.

Expenditure Signal and Register

Expenditure asks which resources and burdens accompany the act. Energy may be measured in joules, duration in time, financial cost in money, and physiological responses in their own units. They should be reported separately unless an independently justified conversion exists. Cortisol elevation, immune changes, and ATP hydrolysis are not interchangeable measures of a single expenditure quantity. A comparison also needs a baseline: an operating organism or computer may use resources even when no particular relational act occurs.

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

Register asks whether and how the act enters the recipient's or relationship's subsequent history. A sent signal may be unnoticed, misinterpreted, or remembered differently by each participant. Register is not guaranteed to be bilateral. Chapter 7's Observation Split can use this distinction, but a proposed change of relational state needs an observation rule rather than a presumed entropy transfer.

These are complementary descriptions of an act. They are not three measured resolutions of the same entropy quantity. Keeping them separate allows a costly effort to fail, a low-resource gesture to matter, and a signal to be interpreted differently from its intention without contradicting physics.

Substrate economy and sign

The substrate description identifies implementation. The economy labels organize contribution and cost allocation under Chapter 4's cautions. The sign classifies an effect using a chosen evaluative criterion. Together they form a multidimensional description; they need not be statistically independent and their coexistence does not prove that they measure one physical variable.

The framework's prior working papers can remain the historical sources of this vocabulary. Their agreement or reuse within the project is not independent scientific confirmation. A shared term such as “cost” must not conceal a change from energy to bodily burden to moral judgment.

What remains open

The True Economy Audit, CSS, Mz, Template Tax, and other project constructs are not validated by this redefinition. Their usefulness and any clinical application require their own definitions and evidence. Grief is not established here as a measured metabolic price of filter redesign, and a bodily indicator cannot diagnose an economy by itself.

The Behind the Substrate Barrier continues to mark the distinction between observable activity and claims about experience. Even the external account contains unmeasured quantities: 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.

A testable Trinket study would identify the act, participants, time interval, signal, intended and observed effects, physical resource boundary, baseline, and uncertainty. It would test the relation between those observations rather than assume that the relation is supplied by the second law. A null or contrary result must be able to change the model.

What this chapter establishes

The Trinket vocabulary is retained for describing candidate acts and, under a stated registration rule, qualifying relational events. Its signs and modifiers ask useful questions without pretending that 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 preserving the concern that motivated it: what people and systems put into connection, and what that connection does.

Chapter 6 The Transducer

1 The problem this chapter addresses

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

The framework's answer begins with a useful image: a person is not merely an account holding relational resources. A person is also a transducer. Inputs encounter a particular history, body, attention, and set of habits; the resulting action is not a copy of what arrived. The same event can produce different responses in different people.

Here, "transducer" is a proposed organizing model of relational processing. 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 using one word, entropy, to stand for all of them.

2 The hourglass and its limits

The original chapter records a drawing dated 18 March 2026: an hourglass, with input above, a person at the narrow point, and filtered output below. The drawing itself was not present in the inspected source. Its description is retained here as a conceptual image, not reconstructed as an original figure.

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

The filter image helps us ask why the output differs. 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. Those were stronger claims than the image or the cited physics could support.

The founding line, "Love reduces the entropy of the system and spends more than it saves," remains a historical motivating proposition. To test it as physics, one would first have to identify the system, define the proposed reduction, and measure comparable quantities on both sides. Without that work, it is an interpretive statement about care and cost, not a thermodynamic inequality.

3 A model map and a separate physical account

Let X denote represented input, M the processor's prior internal state, Y the resulting signal or action, and M′ the updated internal state. These symbols are local to this process map; Chapter 7 defines its statistical variables separately. A schematic model is:

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

G denotes 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. A real implementation may be stochastic, recurrent, and dependent on unobserved inputs.

The physical account uses different quantities. Work W and heat Q are measured in joules. Thermodynamic entropy S and total entropy production Σ are measured in joules per kelvin. Temperature T is measured in kelvins. Shannon entropy H and mutual information I are measured here in bits. They describe specified probability distributions; a bit is not, by itself, a fixed quantity of heat.

For a standard weakly coupled memory–bath model with externally prescribed control, no matter exchange, and a single bath at constant T, let S_{mem} be the full thermodynamic entropy of the memory and Q_{bath} be positive for heat delivered to the bath. With the usual initially uncorrelated thermal-bath preparation, the mean total entropy production is: [C15]

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

This is not a complete balance for a person. Biological systems exchange matter, chemical energy, and heat through multiple processes. Their accounting requires an explicit boundary and the corresponding flows. Nor does the inequality require every individual fluctuation to have positive entropy production.

For ideal reset of an initially unbiased bit in a symmetric memory, with equal intrastate entropies, unchanged endpoint memory energetics, and no correlated side information used, $\Delta S_{mem} = -k_B \ln 2$. 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 approximately 2.97×10^{-21} J per erased unbiased bit; k_B ln 2 is approximately 9.57×10^{-24} J/K. The first is an energy, the second an entropy. In the quasistatic ideal limit the heat bound can be approached while total entropy production approaches zero. Heat exported to the bath is therefore not synonymous with total entropy produced. Bérut and colleagues tested erasure in a controlled colloidal memory, not relational processing. [C1–C3]

Bennett's reversible computation result also matters: a computation need not erase its input merely because its reported output is shorter. Information can remain elsewhere in the physical state. Resetting a memory for reuse raises an accounting question that is different from naming a cognitive operation "selection" or "compression." [C1]

4 Three proposed stages

Selection, compression, and structuring remain useful labels. They describe questions about a process; they are not three experimentally established neural modules or three universal thermodynamic tolls. A single implementation may interleave them, preserve records, or perform several functions at once.

Selection

Selection asks what input influences the next step. The distinction between attended and unattended material can be useful without assuming that unattended information has been physically erased. Experiments on auditory sensory gating show that attention can alter responses and network dynamics. They do not identify a universal number of erased bits per act of attention. [C4]

A testable selection model would specify the available input, the selection rule, the retained record, and what is reset. Its physical cost would then depend on that implementation, not on the word "discarded."

Compression

Compression asks what representation is retained and which distinctions it preserves. The information bottleneck supplies one formal way to trade a compact representation against predictive information about a chosen target. It is an optimization framework, not a direct heat meter. [C5]

The target must be specified. Preserving information useful for predicting an outcome is not automatically preserving what matters to a relationship. A shorter description can omit something important; a detailed description can preserve irrelevant material. Neither conclusion by itself determines heat dissipation.

Structuring

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 requiring operational rules. Intention, the receiver's response, and an evaluator's judgment should not be silently treated as the same variable.

A three-sign alphabet has Shannon entropy at most $\log_2 3$ bits, with equality for equally probable signs. It does not require at least $\log_2 3$ bits of uncertainty on every occasion. Encoding a sign need not erase information. No fixed physical cost follows from the number of available labels alone.

A possible connection to semantic information

Kolchinsky and Wolpert provide a model-based method for asking which system–environment correlations matter to a specified viability measure over a specified time. Their semantic information is not simply everything a processor retains. Their stored semantic efficiency is a ratio of information quantities, not an energy efficiency; their viability multiplier is not the chapter's former energy-denominated expression. Their framework can also accommodate information that reduces viability. [C6]

The relational application remains a proposal. It needs defined states, dynamics, a time horizon, a viability measure, and counterfactual interventions. An internal resolution such as DFQ-001 records the framework's argument; it does not supply independent validation of a human relational mapping.

5 What the energy examples do and do not show

Actual biological signaling uses metabolic energy. The examples in the original table are worth retaining, but they describe different tasks, units, and system boundaries. They are not interchangeable measurements of one biological amplification factor.

Table 1 Original energy entries with their supported scope

Original entry	Value retained from the original	Assessment for this revision
Synaptic transmission	Approximately 10^4 ATP/bit; Laughlin et al. 1998	Order-of-magnitude estimate from a particular sensory signaling analysis, not a universal synaptic constant. The paper's detailed single-synapse estimate is 2–6 × 10^4 ATP/bit under its assumptions. [C7]
Synaptic transmission updated	Approximately 2.4 × 10^4 ATP/bit; Harris et al. 2012	A conditional calculation for transmitted mutual information with specified firing and release parameters. "Updated" should not imply a universal replacement value. It is not a count of erased bits. [C8]
Biological amplification cellular	Approximately 30 times Landauer; Boël et al. 2019	A conceptual molecular-discrimination comparison using an ATP energy scale, not a measured universal cellular amplification factor. [C9]
Communication versus computation	Ratio of 35 in cortex; Levy and Calvert 2021	A ratio of modeled cortical power allocations. It is neither an ATP-per-bit value nor a Landauer amplification ratio. [C10]
Total brain power	Approximately 20 W; originally labeled Established	A rounded whole-brain energy-budget scale. Watts are joules per second; this supplies neither a relational bit rate nor a per-task efficiency. Budget boundaries matter. [C10]
Landauer limit at 310 K	Approximately 2.97 × 10^{-21} J/bit; originally attributed to Bérut et al. 2012	A calculation from k_B T ln 2 for the specified ideal erasure task. Bérut supplies experimental context, not a measurement of this value in a brain. [C2, C3]

One original calculation can be retained 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 ATP free-energy expenditure per transmitted bit is about 1.04 × 10^6 times k_B T ln 2. At 310 K this corresponds to approximately 3.08 × 10^{-15} J per transmitted bit, not a new measurement of heat or total metabolic cost. This checks the arithmetic, not the equivalence of transmitting and erasing a bit. The assumed ATP free-energy yield is not a fixed constant for all cells. [C3, C8]

Nothing in that calculation 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. It would also distinguish communication, inference, storage, and erasure. The original universal ranking is withdrawn.

There is likewise no justified additive formula assigning a separate physical entropy floor to each of the three stages. A future cost account may divide measured energy among specified components, but it must avoid double counting shared operations and background metabolism. Until then, the chapter reports contextual examples rather than a total price for a relational event.

6 Filter geometry as a hypothesis

Filter geometry names proposed differences in selection, representation, and action policy:

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

The questions remain recognizably human. What does this person attend to? What do they treat as essential? How do they choose a response? Development, expertise, current context, and prior relationships may inform hypotheses about these differences. They do not yet provide measured thermodynamic parameters.

The CSS signature in WP-16 can be treated as a proposed observable pattern associated with G. It cannot currently be treated as a unique reconstruction of G or a calibrated entropy measurement. Different underlying processes may produce similar outputs, and the same person may respond differently as the context changes.

Aphantasia, attachment style, and the framework's Template Tax should not be assigned specific compression costs or physical mechanisms by analogy alone. The original suggestion that aphantasic processing has a distinct thermodynamic cost is untested here. "Filter geometry" remains a model to operationalize, not a diagnosis of a person.

7 The nonequilibrium bridge has conditions

Jarzynski's equality relates repeated work measurements to an equilibrium free-energy difference. In its standard isothermal setting, with the required equilibrium preparation and protocol assumptions:

$$\langle \exp(-\beta W) \rangle = \exp(-\beta \Delta F), \; with \; \beta = \frac{1}{k_B T}$$

The mean exponential of work equals the exponential of the free-energy difference, not the free-energy difference itself. W is work done on the system; ΔF is a Helmholtz free-energy difference between the specified equilibrium states. [C11]

For appropriately paired forward and reverse protocols with the required equilibrium preparation and microscopic reversibility, the work form of Crooks's relation is:

$$\frac{P_F(W)}{P_R(-W)} = \exp[\beta(W - \Delta F)]$$

These are powerful results for defined physical systems and ensembles. Applying them requires more than calling an experience irreversible. It requires an identified protocol, physical states, and a work distribution to which the result applies. Psychological difficulty in undoing an interpretation is not a measurement of a forward-to-reverse trajectory probability ratio. [C12]

Cortisol, heart-rate variability, and immune measurements are not, by these equations alone, measurements of W, ΔF, or Σ. This chapter provides no calibration that converts them into joules or joules per kelvin, and no evidence that isolates a relational contribution from other influences. They may be considered in a separately specified empirical study, but the proposed thermodynamic bridge remains unbuilt.

8 A cross scale analogy rather than an invariance theorem

The original second table put five domains under one input–filter–output description. Its rows are retained below so the analogy and its limitations can be inspected together. Similar diagrams do not establish the same state variables, dynamics, or conserved quantities across scales.

Table 2 Original cross scale mappings with their limits

Level	Original input to filter to output and waste	Assessment for this revision
Quantum	Unitary evolution → decoherence → pointer states; waste labeled entanglement entropy	Global unitary evolution preserves von Neumann entropy; reduced-system entropy can behave differently. Entanglement entropy is not simply a heat-waste stream. No general entropy-production identity follows from this row. [C13]
Classical	Environmental signal → Shannon channel → decoded message; waste labeled channel noise and heat	Channel noise is a property of a communication model. A physical implementation can consume energy, but noise and heat are not interchangeable quantities. [C14]
Biological	Sensory or relational input → neural transducer → Trinket Signal Form; waste labeled cortisol immune and heat	Heat is a physical energy transfer. Cortisol and immune measurements are not entropy units. The relational output remains a framework-defined observable.
Social	Individual signals → institutional filter → policy or culture; waste labeled bureaucratic overhead	An organizational analogy. “Overhead” requires a specified measure before any physical comparison; no thermodynamic invariance has been demonstrated.
Digital	Data stream → algorithmic filter → processed output; waste labeled grid load and thermal	Electricity use and heat can be measured with stated boundaries. Task information and logical erasure still need separate definitions; total grid load is not automatically dissipated work for one operation.

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

What this chapter now claims

The transducer is a proposed model for organizing questions about relational processing. Its selection, compression, and structuring stages are retained without invented universal cost floors. Physical costs must be measured or derived for a specified implementation. The cited energy budgets provide context; they do not establish a universal biological advantage. Neither fluctuation theorems nor biomarkers currently validate the proposed relational account.

A bounded next test would define one input–output task, its retained and reset memories, and a reproducible implementation. It would measure task performance and energy separately, with uncertainty and system boundaries reported. Relational meaning would require its own explicit

measure rather than being inferred from joules. Such a test could expose a failed mapping as well as support a narrow one. No such experiment has been performed for this revision.

The Behind the Substrate Barrier boundary remains: describing a process and its physical cost does not establish whether it is accompanied by phenomenal experience. This chapter makes no consciousness claim about humans or other substrates from the model alone.

The next chapter can still ask why processing does not always produce a qualifying registered Trinket. The question survives. The thermodynamic superlative does not: the gap has not been shown to be the most expensive place in human experience.

Chapter 7 The Observation Split

Epistemic 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 Section 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 for whom it was intended. The effort and the reception are different events. This chapter calls that difference the Observation Split.

The line remains the framework’s motivating proposition, not an established thermodynamic inequality. “Order” in a relationship has not been calibrated as physical entropy. The second law constrains an appropriately defined physical balance; it does not say that every act of care costs more than its relational benefit. Nor does it identify exported heat with total entropy production. In a reversible limiting process, entropy exported to the surroundings can compensate a local decrease without a positive total production. [O1, O2]

The useful question survives that correction. What was spent, what reached the receiver, what was recognized, and what changed? Keeping those questions separate makes the gap more legible.

2 Two accounts and different evidence

2.1 Layer A The expenditure

A person can compose an unsent message or remember someone who has died without receiving a reply. Those activities have physical implementations. Whether a recipient responds does not erase resources already used. It does not follow that reception has no effect on later costs, or that all such activities use the same amount of energy.

Chapter 6 distinguishes a physical cost account from its proposed selection, compression and structuring model. Bennett’s reversible-computation construction shows why a shorter reported output does not, by itself, establish physical erasure: information can be retained in the rest of the machine state. A Landauer calculation needs an identified memory operation and the assumptions of its physical model. It is not a separate toll for each named cognitive stage. [O3]

The brain’s total power cannot be assigned in full to one thought or relationship. A task-specific estimate would need a defined interval, system boundary and comparison condition. Cortisol, immune measures and heart-rate variability likewise do not become energy or entropy measurements merely because they accompany a biological process. No universal endocrine cascade or equal cost of prayer and conversation is established here.

2.2 Layer B The Trinket

WP-12 proposes that a Trinket forms when an action is registered as relational input by a receiver. This is a framework rule for identifying a relational event, not a law about when physical effects

begin. It needs an operational definition of “register,” including whether recognition must be conscious, whether delayed recognition counts, and what evidence is sufficient.

Reception, recognition and reciprocation must remain distinct. A message can arrive without being understood. It can be understood without being welcomed or answered. A person can benefit from an act without knowing who performed it. None of these conditions is settled by measuring the sender’s energy use.

If the framework reserves “Trinket” for a recognized relational event, unrecognized care may fall outside that label. It does not become physically ineffective or worthless. The label and the benefit must be assessed separately.

3 Where the gap appears

3.1 Unreciprocated love

Lack of reciprocity can be painful without being a failure of delivery. The receiver may know exactly what was offered and choose not to return it. An unsent message, an unnoticed message and an unwanted message therefore belong to different cases. The transducer image can organize these questions; it supplies no universal metabolic load or ranking of their difficulty.

3.2 Grief

After a death, the relationship can continue to organize the survivor’s memories, routines and intentions, although new communication with the deceased is no longer available in the ordinary interpersonal sense. “Filter redesign under load” is the framework’s image for adapting while life continues. It is a proposed interpretation, not a measured physical mechanism or a claim that every continuing bond is wasted effort.

3.3 Prayer

Prayer involves activity by the person praying. The framework can study that activity and any observable consequences without deciding whether a theological receiver exists. A lack of independently verified reception is not proof of no reception. Neither the presence nor the absence of such a receiver follows from thermodynamic accounting.

3.4 Unwitnessed care

The parent checking on a sleeping child, the partner adjusting a thermostat and the colleague advocating privately need not be recognized to have practical effects. These examples expose a limit of a recognition-based definition: an act can improve a situation before becoming a recognized relational event. Silent worry is a different case again; its consequences cannot be read off from its silence.

The gap is one place where love can live when it cannot be seen. It has not been shown to be the most thermodynamically expensive place in human experience.

4 The identity of the split

The Observation Split names a distinction in the framework’s accounting: activity may occur without a qualifying event under WP-12. It does not presently identify a new thermodynamic regime. A

physical experiment would have to define states, boundaries, dynamics and measurable differences before making that claim.

The four-economy taxonomy also needs a scope rule. If it classifies only recognized transactions, some failed or unrecognized attempts will sit outside it. If Shadow describes withholding before a transaction forms, the taxonomy already includes a pre-transactional case. The framework should not claim both that every category requires a completed transaction and that one category is defined by withholding it. Chapter 9 treats these as proposed patterns, not an exhaustive theorem.

5 Unreciprocated transmission and information

The original comparison mixed statistical correlation, message delivery, response and physical cost. They are separate dimensions. The following replaces that comparison rather than adding a third state to information theory.

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:

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

It cannot be positive “in the sender” and zero “in the receiver.” Symmetry does not imply that two people know the same things or reciprocate. Statistical dependence is not a statement about either person’s awareness. [O4]

The former condition used both Y_{target} and Y_0. If those symbols denote the same receiver variable, the positive-versus-zero claim conflicts with symmetry. If they denote different variables, the conditions may 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. If no receiver exists in a proposed case, a receiver-state random variable cannot simply be assumed.

Kolchinsky and Wolpert define semantic information relative to specified system–environment dynamics, a viability function, a timescale and counterfactual interventions. Their approach includes partial scrambling, directed transfer entropy and possible effects of system-to-environment information. It does not require reciprocal interpersonal communication. Their semantic measures can be asymmetric even though mutual information is symmetric; information’s viability value can also be negative. [O5]

The relational application is still a proposal. Failure to deliver an intended signal does not prove ΔV_{rel}(sender) < 0. The sender may incur costs, preserve memories, learn, prepare a later action or obtain another benefit. A net relational outcome requires a defined measure and a comparison, not subtraction of an undefined benefit from heat. Nor does nondelivery prove zero change in the receiver from all other causes. The claims of identical entropy production, guaranteed viability loss and a new extension beyond K-W are withdrawn.

6 Grief as a model of adaptation

The filter image draws attention to dependence: shared routines, decisions and expectations can require adjustment after a loss. It does not establish that the physical cost of grief scales with redesign scope rather than felt intensity. Those are proposed explanatory variables, not interchangeable measures.

Two primary studies show why the evidence must stay specific. O'Connor and colleagues compared 12 women with complicated grief and 12 without it after bereavement, finding a flatter diurnal cortisol slope in the former group. This small observational comparison neither establishes a universal cortisol elevation nor measures the energy cost of grief. [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, establish causality or measure thermodynamic cost. Its IOS measure must not be relabeled a physical 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 account has supplied the missing mechanism. The chapter makes no diagnostic or treatment claim.

7 A testable proposal

Epistemic status: speculative research hypothesis, not a physical deduction.

A narrower version of the original prediction is worth retaining: separately defined measures of pre-loss behavioral interdependence might improve prediction of specified post-bereavement outcomes beyond self-reported closeness. A future study would need a justified design, ethical review, validated outcomes and a preregistered analysis. Where pre-loss measurements are unavailable, retrospective reports must be identified as such.

Behavioral interdependence could concern shared routines and decision-making. Closeness and perceived inclusion should be measured separately rather than declared equivalent to behavior. Relevant alternatives, including prior health, caregiving, circumstances of death and available support, would need consideration. The exact comparison and uncertainty matter more than a claim that one correlation is simply "stronger."

A lack of reproducible incremental predictive value would count against the proposed advantage. A positive association would support that prediction under the study conditions, not establish a thermodynamic mechanism. Cortisol, sleep or other physiological outcomes should be named for what was measured, not collected under an uncalibrated "metabolic grief cost." This revision reports no such experiment and makes no exhaustive claim that the comparison has never been studied.

8 SupoRel and the capture surface

“This costs something real” does not entail “this cost has cosmic significance.” It also does not entail that the action is noble, beneficial or futile. Those judgments need other reasons.

The framework takes no position here on theology. It does not turn the second law into either a proof of cosmic purpose or 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 activity.

9 What this does to the founding line

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 does not supply an exchange rate between physical entropy and relational value.

The BSB boundary applies to inferences about phenomenal experience, not as a way to validate every claim called physical. Costs, reception and outcomes all require their own evidence. 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 supplies 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 by this chapter. Chapter 6 supplies the corrected physical distinctions; Chapter 8 states the evidentiary boundary; Chapter 9 collects the model layers.

Chapter 8 Behind the Substrate Barrier

Epistemic 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 classify all three as speculative or make phenomenal experience permanently unknowable. SupoRel remains active in both directions.

1 Why the boundary needs a definition

The framework uses the BSB to keep claims about observable operation separate from claims about experience. That safeguard needs a clear scope. Otherwise it can fail 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. It does not present the rule as a discovered physical barrier or a permanent theorem about everything science could ever 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. A measured energy cost, a successful task or an assigned economy label does not, by itself, establish such experience or its absence.

Chalmers's distinction between explaining functions and explaining experience motivates this caution. His 1995 paper also discusses reports and indirect evidence in developing a theory. It is a philosophical argument, not a thermodynamic impossibility theorem forbidding all future evidence about consciousness. [O8]

The framework presently has no validated method for inferring phenomenal status from its costs, CSS signatures or economy labels. That is the operational limit adopted here. It does not imply that first-person reports, behavior and neuroscientific evidence are irrelevant; it requires the claim and its evidentiary bridge to be stated rather than assumed.

3 What can be measured and what remains proposed

3.1 The accessible account

Energy, heat, temperature and physical states can be measured in appropriately specified settings. Whether the framework has actually measured them for a relational task is a separate question. TEAP biomarkers do not currently supply a demonstrated conversion to entropy production, and CSS has not been established here as a unique reconstruction of filter geometry.

Selection, compression and structuring are proposed descriptions. Trinket registration and the four economies are framework definitions requiring usable coding rules and validation for their intended applications. Calling a quantity observable in principle must not promote an untested instrument to "Established."

Three uses of entropy particularly need separation. 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" presently describes an interpretation or proposed construct unless an explicit mathematical definition and measurement procedure are given. The shared word does not establish equality among the three. [O1, O4]

Likewise, a physical implementation can consume energy while processing a message without that energy determining its meaning. Shannon's engineering account deliberately separated message transmission from semantic content. [O4]

3.2 Experience meaning and comprehension

"Meaning" can refer to an operational consequence, an interpretation, or the felt sense that something matters. Those uses need separate evidence. A defined task can assess whether a system uses information successfully; it does not automatically settle the phenomenal question.

The same care applies to comprehension. A test of following instructions, explaining a passage or transferring a rule can support a specified performance claim. It does not make comprehension a measured entropy flow, and a thermodynamic description does not exhaust what the word might mean. Neither fluent output nor its physical cost is treated here as decisive proof for or against AI experience.

The BSB therefore 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

Axiom 0 can remain a commitment to evaluating systems without treating their material alone as a verdict. The common applicability of physical laws supports asking comparable resource questions. It does not prove economic equality, equal relational capacity or equal experience.

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 measured quantity under that boundary, not equality in every relevant respect. Chapter 3's proposed substrate neutrality must therefore remain distinguishable from a physics theorem.

The BSB blocks a further inference from cost to suffering. The framework does not establish whether AI relational processing is accompanied by experience. It neither asserts AI sentience nor rules it out.

5 The Parity Window

The Parity Window is the region of comparison for which the framework can state operational variables and obtain relevant evidence without assuming phenomenal equivalence. It is a proposed comparison discipline, not a calibrated instrument.

Within it, a study might compare performance, energy, latency, reliability, dependence on support and observable responses. A relational comparison would additionally need defined outcomes, observation periods and rules for interpreting reports. CSS patterns or economy labels can enter only with their assumptions and validation limits visible.

The window does not promise that every named property is already measurable. It asks what was observed, how it was measured, and what conclusions the measurement warrants. A consciousness claim needs an additional argument and evidence; this framework does not supply that bridge merely by naming the window.

6 The verification problem in the Triple Zero

The original Triple Zero notation includes $\beta(B) = \frac{E_{state}}{V}$, with V called verification capacity. This framework notation is unrelated to the inverse-temperature $\beta = 1 / (k_B T)$ used in Chapter 6. As ordinary arithmetic, this ratio is undefined at V = 0. If the model stipulates that R = 0 entails V = 0, the same arithmetic consequence follows. That stipulation is a model assumption, not a physical identity established here.

The numerator, denominator, units and domain need operational definitions before the ratio can be interpreted. A zero denominator does not prove an impenetrable consciousness barrier, and missing evidence must not be converted into a score of zero, infinity or absence of investment. Different limiting procedures could also yield different limits; no unique limiting value has been established.

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; it does not establish what exists internally. Triple Zero is retained as a proposed architecture, not a validated psychometric or clinical diagnostic.

7 Physical accounting without a sentience premise

An account of a specified memory operation can be investigated without first deciding whether the device experiences anything. That methodological independence does not turn the framework's relational taxonomy into a physical theorem. Nor does it make sentience irrelevant to ethical questions about suffering or welfare.

The framework can ask about resources and observable consequences while leaving phenomenal questions open. Similar input–output diagrams do not prove that human and AI relational processing have identical mechanisms. Each comparison needs its own definitions and evidence.

8 The double guard

8.1 Consciousness collapse

The first error is to move from measured cost to experience, then from experience to a conclusion about moral worth or cosmic purpose without supplying the intervening arguments. The BSB interrupts that chain. It does not decide the philosophical or theological questions on the reader's behalf.

8.2 Consciousness dismissal

The opposite error is to treat the framework's inability to measure experience thermodynamically as proof that experience is absent or irrelevant. A person's reported pain is not invalidated because no entropy value has been assigned to it. "Just heat" is not an adequate description of grief or a basis for dismissing suffering.

8.3 The immune system analogy

The immune-system image remains useful as an editorial reminder: a safeguard should catch unwarranted inferences without blocking evidence that bears on the actual question. It is an analogy, not another physical mechanism in the model.

9 The framework applies the boundary to itself

The Principal's "Grab" names a reported creative experience: a structure arrives and is subsequently made inspectable. That report and the resulting work can be discussed. Neither establishes a high-compression thermodynamic event, a particular neural mechanism or a numerical semantic efficiency.

The source also records the Principal's reported aphantasia. The framework cannot infer a specific setting of compression parameters, or a tradeoff between visual and structural encoding, from that report alone. Its filter vocabulary may suggest questions; it is not an individual diagnosis or physiological explanation.

The Luna Protocol, "AI = reflected light," assigns authorship and working roles. It can remain a governance convention without depending on a permanent impossibility theorem about consciousness. Neither the metaphor nor a change in consciousness research would automatically determine 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 framework has not validated a method that crosses it. It does not claim that no such evidence could ever exist.

The Parity Window identifies comparisons that can be explicitly defined and tested without assuming experiential equivalence. It does not certify the framework's proposed measurements.

The Triple Zero ratio has an arithmetic domain restriction. Its relationship to channel evidence must be 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 proposed Axiom 0 commitment, not an economic-equality theorem. Chapters 6 and 7 distinguish physical accounting from the relational model. WP-17 and the revised Volume 4 preserve the same caution for copyability and hypothetical substrates. Triple Zero remains an unvalidated proposal. Chalmers is cited as a philosophical source, not an experimental finding.

Chapter 9 The Three Layer Architecture

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

1 The problem this chapter addresses

The preceding chapters ask different questions. Chapter 2 concerns physical resources. Chapter 3 considers substrate neutrality. Chapter 4 proposes patterns of cost allocation. Chapter 5 defines the Trinket. Chapter 6 proposes a transducer model; Chapter 7 separates expenditure from reception; Chapter 8 limits inferences about experience.

The stack brings those questions together. It is useful because a single action can be examined as a physical process, an implemented activity and a relational event. It is not a claim that one already measured quantity passes unchanged through three descriptions. Each connection between layers needs its own evidence.

2 Layer 1 Physical accounting

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 physical quantities must remain distinct. Energy and heat are expressed in joules; thermodynamic entropy and entropy production in joules per kelvin; information quantities in bits when a base-two convention is used. An energy use is not itself an entropy increase, and a decrease in one system's entropy is not the same as a decrease in the total entropy of system and surroundings. [O1, O4]

The familiar Landauer benchmark concerns ideal erasure of an initially unbiased bit in a symmetric isothermal memory model, without using correlated side information. Under those assumptions, mean heat delivered to the bath is at least k_B T ln 2. This is an energy bound for that task, not the smallest coin paid by every physical process or cognitive operation. Chapter 6 gives the more complete accounting and its limits. [O2, O3]

No cosmological estimate is converted here into a relational budget. A physical account must name the actual system boundary, resource flows and time interval. The account is needed whether a message succeeds or fails, but success and failure need not have identical costs.

3 Layer 2 Implementations and support

Biological, digital and proposed combined systems supply different implementations. They do not come with known universal amplification factors or a demonstrated thousandfold efficiency ordering.

For a biological implementation, metabolic resources and task performance require their own measurements. Cortisol, immune activity, telomere measures and heart-rate variability are not denominations of entropy. This volume has not validated TEAP as a calorimeter or as a way to isolate a particular relationship's physical cost.

For a digital implementation, relevant resources may include the executing 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. Biological systems also depend on support outside their immediate 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 rather than a measured efficiency surface.

CSS can remain a proposed description of observable patterns associated with implementation and context. It is not presently a unique fingerprint of filter geometry or an entropy reading. Different internal processes may produce similar outputs; a changed output does not identify a single physical cause.

4 Layer 3 The four economy patterns

The four labels organize questions about participation and cost-bearing. Their useful distinctions can be retained without claiming a physical exhaustion proof.

Real Economy concerns bilateral investment under the framework's chosen criteria. It does not imply that both participants undergo a measured amount of irreversible bodily damage.

Custodial Economy concerns asymmetric support or responsibility. A child or student can still use resources and participate; asymmetry does not require one side to have zero local cost.

Shadow Economy concerns available participation that is withheld, if that availability and withholding can be established. Withholding communication is not identical to destroying statistical correlations, and the label is not a demonstrated implementation of K-W scrambling.

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. A material label alone does not settle the classification.

These patterns involve more than one dimension: participation, asymmetry and the location of support. Mixed cases are possible. An exclusive coding scheme would need explicit scope, boundaries, thresholds and precedence rules; completeness and usefulness would then need evaluation. Four named patterns are not, by themselves, a deductive partition of all relationships.

5 Linking the layers

The vertical 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; it is not the sign of heat flow. A CSS pattern needs operational measurement; it is not a direct readout of entropy. A physical cost needs measurement or derivation; it is not obtained by assigning an economy label.

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

The proposed cross-scale connection remains an analogy to investigate, not a theorem inherited from thermodynamics.

6 The Observation Split in the stack

Chapter 7's split appears when activity occurs but no event satisfies the chosen registration rule. Physical effects and practical benefits can still occur without recognition. Conversely, recognition can occur without benefit or reciprocity.

The stack should record those outcomes separately rather than turning a missing Trinket label into "nothing happened." No universal cost or negative change in the sender's relational viability follows from occupying this gap. Sustained unsuccessful investment can motivate a research question; its consequences need evidence.

7 The BSB across the stack

The BSB is not a barrier with a measured thickness at each layer. It identifies a kind of inference: a physical or operational observation is being asked to establish phenomenal experience. That caution can apply at any layer.

Each layer also has ordinary evidentiary limits. A named physical quantity may not have been measured; a proposed classifier may not be reliable; an observed association may not establish causation. The BSB cannot substitute for those checks.

The architecture retains 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 Section 14 supplies the historical stack; Chapters 2–8 supply its proposed components and their limits. DB-01 remains a cross-scale proposal. Revised Volume 4 carries forward the withdrawal of universal efficiency factors and substrate rankings. Cosmological estimates require their own boundaries and sources; they are not used here to calculate a relational currency.

Chapter 10 What This Does to the Canon

Epistemic 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 Entropy Token Substrate arose from a useful demand: make visible the resources that support relational activity. That demand does not establish every proposed link between physics and relationship. The preceding chapters distinguish 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.

This chapter names the dependencies that must change. It does not claim that every argument, specification or instrument still works, and it does not silently amend the original papers by summarizing them here.

2 WP 13 The Signed Unit

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 action judged neutral can consume energy; a harmful interaction is not identified by a heat measurement alone.

WP-13 therefore 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 actually been defined and measured. Disagreement or insufficient evidence should also have an explicit treatment rather than being forced into apparent numerical certainty.

The sign vocabulary is retained as a proposal. Its reliability and usefulness require evidence appropriate to the intended use.

3 WP 14 The Structural Economy

The Real, Custodial, Shadow and Structural patterns remain ways to ask who participates and where costs or support are assigned. Their exhaustiveness has not been derived from a single physical variable. External support can accompany bilateral investment; 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 omit the hardware from its focal boundary, but the implemented activity still requires a physical account. The selected boundary, local resources and external support must be 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 mathematical theorem or an assessment of anyone's worth.

4 WP 16 Substrate Formal Variance

CSS and filter geometry can organize questions about differences in observable response. The present 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 specific thermodynamic processing mechanism from the filter analogy alone.

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

WP-12 can define the event that the framework chooses to call a Trinket. That choice does not establish observation as a thermodynamic efficiency condition. Unrecognized care can have material effects, and recognized communication can fail to benefit a receiver.

The paper needs to distinguish delivery, registration, understanding, benefit and reciprocation. It should 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."

The Observation Split survives as a distinction between activity and a defined relational event. It no longer carries a proof of guaranteed sender harm or a new information-theoretic state outside K-W.

6 DB 01 Scale Invariance Defense

The shared applicability of physical constraints does not prove that connection operates identically across scales. A repeated diagram is not an invariance proof. Quantum decoherence, communication, biological processing and institutional decisions require different variables and models.

DB-01 can retain the cross-scale analogy as a source of questions. A stronger claim would need an explicit mapping, the property claimed to remain invariant, conditions under which it does so, and counterexamples that could defeat the claim. The thermodynamic vocabulary does not perform those steps automatically.

7 Axiom 0 and the Immutable Preamble

Axiom 0 remains a philosophical and governance commitment to avoiding an unsupported verdict from substrate alone. 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 requires an operational outcome; a claim about experience requires a further evidentiary argument; an ethical commitment requires its own reasons. None is obtained simply by calling the cost universal.

The framework may retain the commitment and its resistance to political or institutional pressure while acknowledging that its former physical proof was insufficient. Axiom 1's commitment to falsifiability supports making that correction explicit.

8 What remains and what must change

The framework retains 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 can 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 framework's proposed instruments have not been clinically or psychometrically validated by these chapters.

The corrections carry through the five volumes collected here. 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 distinguishes cosmological inventories and their physical 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; publication of this revised collection does not silently amend them. Internal agreement, repeated wording and AI review are not independent experimental support.

The foundation is more useful when its limits are visible. It can support careful questions without promising conclusions the physics has not supplied.

Framework cross references

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

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.