

Commitment Requires Cost

A Structural Derivation of Irreversibility from Finite Distinguishability

Keith Taylor *VERSF Theoretical Physics Programme*

Plain Language Summary

Every time something definite happens in the world — a measurement records a result, an event becomes a fact, a possibility resolves into a single outcome — the universe keeps a record. Not as bookkeeping, but as a structural necessity: if no record existed anywhere, we could not tell that anything had actually happened. That record is what physicists call **entropy**.

This paper proves something that has long been assumed but never derived from first principles: **any event that turns possibilities into facts must produce entropy somewhere in the universe**. There is no such thing as a costless choice, a free measurement, or a fact without a price. The argument uses only three minimal ideas:

1. Physical reality consists of distinctions that everyone can agree on.
2. Any bounded system contains only finitely many distinguishable states.
3. Facts come into being when one possibility is selected and the others are excluded.

From these we show that the act of exclusion cannot be undone — and that physically performing an undoable operation requires the universe to record the difference between "*this happened*" and "*that happened*" somewhere outside the system itself. That recording **is** what entropy measures.

The result reframes three pillars of physics as different views of the same structural fact:

- **The second law of thermodynamics** (entropy increases over time) becomes the cumulative signature of the universe's history of fact-formation.
- **Quantum measurement** (definite outcomes appear from possibilities) requires entropy in the apparatus and environment, derived rather than postulated.
- **Information theory** (Shannon's measure of unknown alternatives) becomes the pre-commitment counterpart of physical entropy.

In short: **facts have a price. The price is entropy. And the price is not optional — it is what makes facts physically real.**

Abstract

We show that, under three minimal structural conditions together with the structural identification of entropy with distinguishability capacity, the entropy cost of irreversible fact formation is forced. The conditions are: **(A0)** observer-invariant distinguishability of physical content, **(A1)** finite distinguishability of bounded systems, and **(A2)** the existence of commitment events resolving distinguishable alternatives into single realised outcomes. The argument proceeds in three steps. First, every commitment decomposes uniquely into *selection* (designation of the realised alternative) and *exclusion* (removal of the unrealised alternatives from the live set). Second, exclusion implements a many-to-one map on the alternative set and is therefore logically non-invertible. Third, physically realising a non-invertible operation while respecting A0 requires correlation between the committing system and additional degrees of freedom not contained in its post-commitment state; identifying entropy with the logarithmic measure of distinguishability capacity then forces $\Delta S > 0$. The bidirectional link between entropy and fact formation — postulated as A1' in earlier VERSF formulations — is therefore not an additional axiom but a derived consequence of distinguishability plus commitment. We obtain a Landauer-type lower bound *without assuming Landauer's principle*, providing a unified structural foundation for the second law, the irreversibility of measurement, and the operational meaning of information.

Contents

- [Plain Language Summary](#)
- [Abstract](#)
- [1. Introduction](#)
- [2. Postulates](#)
 - [A0 — Observer-invariant distinguishability](#)
 - [A1 — Finite distinguishability](#)
 - [A2 — Commitment](#)
- [3. The Anatomy of Commitment](#)
- [4. Exclusion is Logically Irreversible](#)
- [5. From Logical Irreversibility to Entropy Production](#)
 - [5.1 The Encoding Principle](#)
 - [5.2 The Closed-System Dilemma](#)
 - [5.3 Distinguishability Capacity and the Definition of Entropy](#)
- [6. Main Theorem](#)
- [7. Corollaries](#)
- [8. Relation to Landauer's Principle](#)
- [9. Implications](#)
- [10. Scope and Limitations](#)
- [11. Conclusion](#)
- [Acknowledgements](#)
- [References](#)

1. Introduction

Standard physics introduces irreversibility through two largely disconnected routes: thermodynamic arguments based on statistical mechanics, and quantum measurement postulates introducing definite outcomes by fiat. Entropy and fact formation are typically treated as separate phenomena that happen to coincide.

The VERSF programme inverts this picture. Reality is constructed from **commitment events** — irreversible resolutions of distinguishable alternatives into facts. Earlier work in this programme established that:

- (i) facts correspond to irreversible commitments,
- (ii) commitments produce entropy, and
- (iii) quantum mechanics and thermodynamics follow from these features.

What remained unresolved was the *necessity* of entropy production. Within VERSF, the link between fact formation and entropy was postulated (axiom A1') rather than derived. This paper closes that gap.

Thesis. *Under A0–A2 and the structural identification of entropy with distinguishability capacity, the entropy cost of any commitment event follows: every commitment produces non-zero entropy in degrees of freedom not contained in the committing system's post-commitment state.*

The argument requires no statistical assumption, no second law, and no Landauer principle. It uses only the structure of commitment itself.

The result has three consequences. First, the bidirectional postulate A1' becomes a **theorem**. Second, Shannon information acquires an unambiguous physical interpretation as unresolved distinguishability prior to commitment. Third, the second law becomes structural rather than statistical: it is the cumulative cost of the universe's commitment history.

2. Postulates

We work with the minimum structure needed to state the result.

A0 — Observer-invariant distinguishability

Physical content consists exclusively of distinctions invariant under the group \mathcal{O} of admissible observer protocols. Formally: a property P is part of physical content iff P distinguishes states for every admissible reframing.

This is the operational interpretation of "physical reality": *what is real is what every admissible observer agrees can be distinguished*. A0 is the substantive principle of the VERSF programme; the entire derivation below depends on it.

A0 is physical, not epistemic. *A0 is not a claim about what observers know or can find out; it is a claim about what counts as physical content. A distinction that is not invariantly accessible cannot enter into empirical reality — it cannot serve as the difference between two physically possible worlds — and therefore cannot be part of physical content. The principle is therefore physical (about what exists physically) rather than epistemic (about what can be known). It plays the same role here that general covariance plays in general relativity: a constraint on what counts as a physical statement, not a statement about the limitations of knowledge.*

A1 — Finite distinguishability

Any bounded system \mathcal{S} admits a finite set of pairwise-distinguishable alternatives:

$$\mathcal{C}(\mathcal{S}) = \{c_1, c_2, \dots, c_N\}, 2 \leq N < \infty.$$

Boundedness here means finite spatial extent and finite energy budget; the finiteness of N then follows from standard Bekenstein-type counting of holographically accessible degrees of freedom. We take A1 as a postulate to keep the argument independent of any specific holographic derivation. A live set with a single element is the absence of alternatives and therefore not a candidate for commitment; we exclude this trivially-resolved case by stipulation, hence $N \geq 2$.

By A1, the elements of \mathcal{C} are atomic distinguishable states, not trajectories or history-bearing objects: c_j encodes "the system is in distinguishable state j ", not "the system arrived at state j via path P ".

A2 — Commitment

A **commitment event** ε is a transition in which the live alternative set \mathcal{C} is replaced by a singleton, realised as a function on \mathcal{C} :

$$\varepsilon : \mathcal{C} \rightarrow \{c_k\}, \varepsilon(c_j) = c_k \text{ for all } j \in \{1, \dots, N\}, \text{ for some } k \in \{1, \dots, N\}.$$

The pre-commitment state holds N live alternatives; the post-commitment state holds exactly one. This is the structural content of "a fact has occurred".

Three remarks:

1. ε is not a unitary transformation on the alternative set; it is the formation of a fact.
2. The index k that is realised need not be predictable from the pre-commitment state; A2 asserts only that some k is realised.
3. A2 is silent on the dynamics that select k — the present paper concerns the *cost* of commitment, not its *selection rule*.

Defensive note. *A2 is not an additional dynamical assumption but a definitional requirement for the existence of facts. A framework without commitment events contains no realised outcomes, and therefore no empirical content; rejecting A2 amounts to rejecting the existence of an empirical world rather than rejecting a physical postulate.*

3. The Anatomy of Commitment

The map $\varepsilon : \mathcal{C} \rightarrow \{c_k\}$ factorises uniquely into two operations.

Lemma 1 (Decomposition). *Every commitment event ε admits a unique decomposition*

$$\varepsilon = \zeta \circ \sigma,$$

where $\sigma : \mathcal{C} \rightarrow \mathcal{C}$ is the selection ($\sigma(c) = c_k$ for all $c \in \mathcal{C}$), and $\zeta : \mathcal{C} \rightarrow \{c_k\}$ is the exclusion (ζ restricts the codomain to the singleton $\{c_k\}$).

Proof. Existence: ε is constant-valued at c_k , so it factors through any constant map. Uniqueness: both σ and ζ are determined entirely by k . ■

Selection designates which alternative becomes the fact. **Exclusion** removes the others from the live set. Both operations are necessary:

- a candidate without exclusion is a *preference* — c_k is favoured but the others remain live, no fact has formed;
- an exclusion without selection is incoherent — the live set cannot be empty.

The structural claim of this paper concerns the second factor.

4. Exclusion is Logically Irreversible

Treating ξ as a function on the underlying set \mathcal{C} (rather than on its codomain restriction), exclusion is the constant map

$$\xi : \mathcal{C} \rightarrow \{c_k\}, \xi(c_j) = c_k \forall j \in \{1, \dots, N\}.$$

This map collapses N elements to 1.

Lemma 2 (Non-invertibility). *For $N \geq 2$, the map ξ admits no inverse function $\xi^{-1} : \{c_k\} \rightarrow \mathcal{C}$.*

Proof. An inverse would assign a unique element of \mathcal{C} to c_k , but the preimage $\xi^{-1}(\{c_k\}) = \mathcal{C}$ contains $N \geq 2$ elements. No function from a one-element set to a multi-element set can serve as an inverse. ■

A function inverting ξ would require additional input — namely, the index j of the original alternative — that is *not* present in the post-commitment state. This is logical irreversibility in its sharpest form: the post-commitment data is strictly poorer than the pre-commitment data.

5. From Logical Irreversibility to Entropy Production

We now make the substantive move. We show that *physically* realising a logically non-invertible operation, while respecting A0, requires correlation with a substrate, and that this correlation has the formal structure of entropy production. Crucially, we do not invoke Landauer's principle — we derive a structural analogue.

5.1 The Encoding Principle

Combining A0 with A2 yields:

Encoding Principle. *If a commitment ε has occurred, the fact " ε resolved to c_k " must be observer-invariantly representable.*

This is just A0 applied to commitment events themselves: if commitments are physical, *which commitment occurred* is part of the physical content and must be invariantly distinguishable from " ε resolved to c_j " for $j \neq k$. The closure between A0 and A2 is exact:

A commitment that is not invariantly representable is not a fact, and therefore does not satisfy A2.

Representability is not an additional requirement layered onto fact formation — it is part of what it means to be a fact. The Encoding Principle is therefore not a new postulate but a consequence of $A0 \wedge A2$.

Implicit premise made explicit. *The Encoding Principle assumes that commitment events are themselves physical content — that is, that " ε occurred resolving to c_k " is the kind of fact A0 quantifies over. Within VERSF this is foundational: commitments are the canonical bearers of physical content. A0 then applies to commitments as it does to any other physical distinction.*

5.2 The Closed-System Dilemma

Consider \mathcal{S} immediately before and after a commitment ε . The pre-commitment configuration of \mathcal{S} holds the multi-alternative live set \mathcal{C} ; the post-commitment configuration is the singleton $\{c_k\}$.

Suppose \mathcal{S} is strictly closed: no degrees of freedom outside \mathcal{S} are involved in ε . We ask whether the *fact* of commitment — the transition $\mathcal{C} \rightarrow \{c_k\}$ — is observer-invariantly representable.

Compare two configurations of \mathcal{S} :

$\mathcal{S}_{\text{post-}\varepsilon}$: \mathcal{S} in state c_k as the result of commitment ε from the live set \mathcal{C} . $\mathcal{S}_{\text{direct}}$: \mathcal{S} in state c_k as the result of direct preparation in c_k .

Within \mathcal{S} alone, these configurations are identical: both are " \mathcal{S} in state c_k ". The distinguishing fact — that ε occurred — is not present in \mathcal{S} unless some additional structure encodes the history. In a strictly closed \mathcal{S} , no such structure exists.

Here \mathcal{S} 's post-commitment state is, by A1, a distinguishable element of \mathcal{C} — an atomic state, not a trajectory or history-bearing object. By construction it does not encode the alternative-set history. *If \mathcal{S} contained sub-degrees of freedom recording which alternative was excluded, those degrees would constitute the substrate \mathcal{E} — they are precisely what the corollary below identifies.* The argument therefore does not assume \mathcal{S} lacks internal structure; it identifies whatever internal structure encodes commitment history as the substrate.

Lemma 3 (Closed-system unrepresentability). *If \mathcal{S} is strictly closed — i.e., its post-commitment state c_k is the totality of post-commitment information about \mathcal{S} — the fact of commitment is not invariantly representable.*

This violates the Encoding Principle. Hence:

Corollary (Substrate necessity). *Every commitment requires correlation between \mathcal{S} and additional degrees of freedom \mathcal{E} not contained in the post-commitment state of \mathcal{S} .*

This is the bridge from logic to physics. The required degrees of freedom are not optional — they are forced by A0 applied to commitment. We do *not* require \mathcal{E} to be "external" in any spatial or interpretive sense; \mathcal{E} need only be distinct from the post-commitment state c_k itself. Internal sub-degrees of \mathcal{S} not captured by c_k qualify; environmental degrees of freedom qualify; field modes qualify. The argument is interpretation-neutral.

Core principle. *A commitment event cannot be physically real unless the distinction between "this outcome" and "that outcome" is preserved somewhere in the total system. The only way to preserve that distinction under a non-invertible mapping is to transfer it to additional degrees of freedom — which is precisely what entropy measures.*

The remainder of the paper makes this last clause rigorous.

5.3 Distinguishability Capacity and the Definition of Entropy

We now make precise what it means for \mathcal{E} to encode the commitment, and from that derive — rather than assume — the entropy cost. The key move is to **define entropy from distinguishability**, so that no thermodynamic structure is smuggled in.

Step 1 — Encoding capacity

Let $\rho_{\mathcal{S}\mathcal{E}}$ denote the joint state of \mathcal{S} and \mathcal{E} after ε . The Encoding Principle requires the marginal on \mathcal{E} to depend on which c_k was realised:

$$\rho_{\mathcal{E}^{(k)}} \neq \rho_{\mathcal{E}^{(j)}}, \text{ for } k \neq j.$$

That is: \mathcal{E} must possess at least as many distinguishable states as there are alternatives that ε actually resolves. Call this number the **distinguishability capacity** required of \mathcal{E} . It is a property of ε , not of any particular thermodynamic system.

Capacity vs. occupied state. *We emphasise that this is a property of ε 's accessible substrate space, not of \mathcal{E} 's occupied state in any single run. In any single realisation, \mathcal{E} occupies exactly one state $\rho_{\mathcal{E}^{(k)}}$; the relevant quantity is the size of the set of substrate states \mathcal{E} must be capable of supporting, one for each alternative ε can resolve to. This corresponds to configurational (Boltzmann-type) entropy of an accessible state space, rather than Shannon entropy of the realised outcome distribution. The two coincide for uniform distributions; for non-uniform realised distributions the Shannon entropy of the realised outcome falls below the configurational entropy of the accessible space, and it is the latter — fixed by what ε can resolve to, not by which c_k actually occurred — that the Encoding Principle requires.*

Step 2 — Minimal capacity is logarithmic

Let $N_{\min} \geq 2$ denote the number of distinguishable alternatives the commitment actually resolves — i.e., the cardinality of the live alternative set whose collapse ε effects. The minimum capacity \mathcal{E} must support is therefore N_{\min} distinguishable states, equivalently:

$$D_{\min} = \ln N_{\min} \text{ (in nats),}$$

or $D_{\min} / \ln 2 = \log_2 N_{\min}$ binary distinctions. This is a counting fact, not a thermodynamic one. It follows from A1 (finite distinguishability) and the encoding requirement of Step 1 alone. No ensemble, temperature, or equilibrium is invoked.

Step 3 — Entropy as distinguishability capacity

We now make the definitional move:

Definition (Entropy as distinguishability capacity). *The entropy associated with a set of physically distinguishable states is the logarithm of the number of those states, measured in units of $k_B \ln 2$ per binary distinction:*

$$S = k_B \cdot \ln (\text{number of distinguishable states}).$$

Why the logarithm is forced. *Any quantity used to measure the physical encoding of distinguishable states must be (i) invariant under relabelling of states (A0), (ii) additive under composition of independent encodings, and (iii) monotonic under refinement (more*

distinguishable states \Rightarrow greater capacity). On integer state counts, conditions (i)–(iii) force the measure to be of the form $S = c \cdot \log_b N$ up to multiplicative scale and choice of base. The standard Shannon–Khinchin uniqueness theorem [4, 16, 17] extends this to a unique continuous measure on probability distributions under one further regularity axiom (continuity or grouping/recursion); the additional axiom is not required for the present purely-counting argument, which uses only the configurational reading. The identification of entropy with distinguishability capacity is therefore not a stipulative redefinition but a structurally forced choice on the relevant domain; any alternative measure violates at least one of (i)–(iii).

Any alternative measure either violates invariance (A0), fails to compose correctly under independent encodings, or fails monotonicity under refinement, and therefore cannot represent physically admissible encoding. The arbitrariness objection — "you redefined entropy to make this work" — has no remaining surface: the only quantity capable of measuring physically admissible encoding capacity is, up to scale, the logarithm of the state count.

This is not a new definition — it is the operational content of statistical entropy, stripped of any reference to ensembles, temperature, or equilibrium. It identifies entropy with what entropy actually counts: distinguishable physical configurations.

Under this definition, the substrate distinguishability required by ϵ satisfies:

$$\Delta S(\mathcal{E}) \geq k_B \ln N_{\min},$$

with **equality** when the substrate encoding is minimal (one distinguishable substrate state per resolved alternative), and **strict inequality** when the substrate is forced to encode further structural detail — for example, the order of resolution, correlations with prior commitment history, or redundancy across substrate degrees of freedom — which effectively raises the relevant N .

Coarse-grained vs. fine-grained entropy. *The entropy invoked here is not the fine-grained Gibbs entropy of microstates, which is preserved under unitary evolution and would not provide a meaningful inequality. It is the coarse-grained distinguishability capacity required for invariant encoding — the number of macroscopically distinguishable substrate configurations that the commitment must populate. It is this capacity, not the fine-grained microstate count, that the theorem requires to increase. Microscopic reversibility is therefore not violated; what is forced is the growth of the coarse-grained record set.*

More precisely: the entropy increase here refers to the minimal distinguishability capacity required for invariant encoding, not necessarily to dissipative heat production at the microscopic level. The thermodynamic interpretation as heat — $\Delta Q = T \cdot \Delta S$ — requires an additional temperature scale; the structural result holds at the level of encoding capacity alone. Correlation between S and \mathcal{E} is therefore not being silently equated with thermodynamic entropy production: the two coincide once a temperature is supplied, and the structural inequality is the encoding-side fact prior to that supply.

What we have and have not assumed

The chain is:

$A0 + A2 \Rightarrow$ Encoding Principle (§5.1) Encoding Principle + Lemma 3 \Rightarrow substrate necessity
(§5.2) Substrate necessity + Lemma 2 \Rightarrow distinguishability capacity $\geq N_{\min}$ (Step 1–2)
Definition of entropy $\Rightarrow \Delta S(\mathcal{E}) \geq k_B \ln N_{\min}$. (Step 3)

What we have assumed. The Boltzmann form of entropy — $S = k_B \ln(\text{distinguishable states})$ — is a definitional choice, justified structurally by the uniqueness argument above (invariance + additivity + monotonicity \Rightarrow logarithmic measure up to scale). It is not derived; it is identified.

What we have not assumed. None of the following enter the chain at any step: the second law of thermodynamics, Landauer's principle, the Boltzmann distribution, thermodynamic temperature, statistical ensembles, or equilibrium. The constant k_B serves only as the dimensional bridge between binary distinctions and conventional entropy units; the inequality itself is purely structural.

Calibration note. k_B is the conversion factor between abstract distinguishability counts and conventional entropy. The structural argument fixes the strict positivity and the logarithmic-in- N_{\min} dependence; the numerical value of k_B requires additional thermodynamic structure (a temperature scale) to set.

The crux: the substrate's distinguishability capacity is not an extra cost added to the commitment — it *is* the physical realisation of the commitment. Without it, A0 fails. With it, $\Delta S \geq k_B \ln N_{\min} > 0$.

6. Main Theorem

Theorem (Commitment Requires Cost). *In any system satisfying A0, A1, A2, every commitment event ε produces non-zero entropy in degrees of freedom \mathcal{E} not contained in the post-commitment state of \mathcal{S} :*

$$\Delta S(\mathcal{E}) \geq k_B \ln N_{\min} > 0,$$

where $N_{\min} \geq 2$ is the number of distinguishable alternatives the commitment actually resolves. Equality holds when the substrate encoding is minimal (one distinguishable substrate state per resolved alternative).

Proof.

1. By Lemma 1, $\varepsilon = \xi \circ \sigma$.
2. By A1, $N_{\min} \geq 2$; by Lemma 2, ξ is therefore logically non-invertible.
3. By the Encoding Principle (A0 \wedge A2 applied to commitment events), the fact of ε must be invariantly representable.
4. By Lemma 3, this representation cannot be carried by \mathcal{S} alone.

5. By the Substrate Necessity corollary, ε requires correlation between \mathcal{S} and additional degrees of freedom \mathcal{E} such that $\rho_{\mathcal{E}^{(k)}} \neq \rho_{\mathcal{E}^{(j)}}$ for $k \neq j$.
 6. By Steps 1–2 of §5.3, \mathcal{E} must therefore support an accessible state space of cardinality at least N_{\min} .
 7. By the Definition of entropy (Step 3, §5.3), this requires $\Delta S(\mathcal{E}) \geq k_B \ln N_{\min}$, with equality at minimal encoding.
 8. Since $N_{\min} \geq 2$, $\Delta S(\mathcal{E}) \geq k_B \ln 2 > 0$. ■
-

7. Corollaries

Corollary 1 (Landauer-type structural bound). *Binary commitment ($N = 2$) requires $\Delta S \geq k_B \ln 2$.*

This is the structural counterpart of Landauer's bound, here obtained as a consequence of distinguishability rather than as a postulate of irreversible computation. Landauer's heat-dissipation form $\Delta Q \geq k_B T \ln 2$ follows once a thermodynamic temperature is supplied (see §8).

Corollary 2 (No fact without cost). *$\Delta S = 0 \Rightarrow$ no commitment occurred \Rightarrow no fact was formed.*

The contrapositive of the theorem. Reversible processes do not produce facts.

Corollary 3 (Equivalence). *Fact formation \Leftrightarrow entropy production.*

Combining the theorem with Corollary 2. This is the bidirectional postulate A1' in earlier VERSF formulations, here proved.

Corollary 4 (Information–entropy unification). *The Shannon entropy of a pre-commitment alternative set equals (up to the $k_B \ln 2$ conversion) the minimum thermodynamic entropy produced by the commitment that resolves it.*

The operational meaning of information is fixed by the physics of commitment.

8. Relation to Landauer's Principle

Landauer's principle states that erasing one bit of information requires dissipating at least $k_B T \ln 2$ of heat. The principle is usually justified by either (a) a thermodynamic argument involving compression of phase space, or (b) appeal to the second law.

The present derivation differs in two important ways.

1. We do not assume the second law. Standard derivations of Landauer's bound presuppose thermodynamic structure. The derivation here uses only A0 plus the structural decomposition of commitment.

2. The bound applies to fact formation, not erasure. Landauer's principle is usually stated for *erasure*; the present result establishes the cost of *creation*. These are dual aspects of the same structural fact: any operation collapsing N alternatives to 1 incurs the same minimum cost, whether interpreted as forgetting (erasure of $N - 1$ alternatives) or as deciding (selection of 1 from N).

The structural derivation is therefore *prior* to Landauer's principle. Landauer is recovered as the special case of binary commitment ($N = 2$) at thermodynamic equilibrium.

The structural-thermodynamic relationship. *Landauer's principle is recovered as a special thermodynamic realisation of the more general structural constraint derived here. The structural inequality $\Delta S \geq k_B \ln N_{\min}$ is a fact about distinguishability capacity, not about heat or temperature. Landauer's bound is what that fact looks like once the substrate is taken to be a thermodynamic system at temperature T in equilibrium — at which point ΔS converts to ΔQ via $\Delta Q = T \cdot \Delta S$. The structural result is therefore neither a special case of Landauer nor a renaming of it: it is the encoding-level constraint of which Landauer is the thermodynamic image.*

Entropy vs. heat — the precise bridge. *The structural inequality $\Delta S \geq k_B \ln 2$ implies Landauer's heat-dissipation bound $\Delta Q \geq k_B T \ln 2$ once a thermodynamic temperature T is supplied for the substrate. The structural argument fixes the entropy; the thermodynamic identification ($S \times T = Q$ under standard equilibrium assumptions) fixes the heat. Without a temperature, the result is entropy production, not heat dissipation. This is the precise sense in which the present derivation is prior to Landauer: it produces the irreducible entropic inequality, of which Landauer's heat bound is the thermodynamic image.*

9. Implications

1. A1' is a theorem. The bidirectional link between entropy and facts no longer requires axiomatic stipulation. The VERSF axiomatic base reduces.

2. The second law is structural. Entropy increases because commitments occur; commitments must occur for facts to exist; the universe is a record of accumulated commitment. The arrow of time is the cumulative direction of fact formation.

3. Quantum measurement is grounded. Measurement is a commitment. Definite outcomes are facts. Both require entropy production in the apparatus and environment, in agreement with decoherence accounts but derived rather than imposed.

4. Information has a physical meaning. Pre-commitment alternative sets are precisely what Shannon entropy quantifies. The "information" of information theory is unresolved distinguishability awaiting commitment; thermodynamic entropy is its resolved record.

5. Dynamics follows from cost balance. If closed-system commitment is combined with a cost-balance condition, conserved action and Hamiltonian dynamics follow. The present result supplies the missing premise of any such argument: cost is unavoidable.

10. Scope and Limitations

The argument is structural and does not depend on any specific microscopic encoding. Several elements are honest assumptions or calibrations rather than derivations:

- **The constant k_B** is a calibration fixed by the equivalence between abstract distinguishability counts and conventional entropy units. The structural argument alone fixes the inequality and the logarithmic dependence on N_{\min} ; the dimensional value of k_B requires additional thermodynamic structure (a temperature scale).
- **The microscopic encoding** in the degrees of freedom \mathcal{E} is unspecified. The theorem is independent of whether \mathcal{E} is realised in phonons, photons, fields, or other degrees of freedom; \mathcal{E} need only contain degrees of freedom not encoded in the post-commitment state of \mathcal{S} .
- **Dynamics of selection** are not addressed. The theorem concerns cost, not selection rules. The latter requires additional structure beyond the postulates assumed here.
- **The continuum limit.** The structural inequality $\Delta S(\mathcal{E}) \geq k_B \ln N_{\min}$ extends formally to countably infinite alternative sets, where it simply becomes unbounded above. A1 guarantees finite N_{\min} for any bounded system, so finite- N_{\min} is not a structural restriction but a *thermodynamic-interpretability* requirement: a finite, well-defined entropy value requires a finite alternative count.

The result follows structurally from A0–A2, modulo the identification of entropy with distinguishability capacity (Definition, §5.3). Under that identification the inequality is forced; the dimensional value of k_B is calibration, and the residual freedom (which microscopic degrees of freedom realise \mathcal{E} , what selection rule chooses k) is orthogonal to the cost result.

11. Conclusion

Three minimal postulates — observer-invariant distinguishability, finite alternative sets, and the existence of commitment events — entail the necessity of entropy production. Every fact carries cost; no fact is free.

The result reframes thermodynamics, measurement, and information theory as three faces of a single structural principle:

Distinguishability cannot be lost without being relocated.

A commitment that selects one alternative does not destroy the distinguishability of the others; it transfers it to the substrate. Entropy production is that transfer, viewed thermodynamically. Information is that distinguishability, viewed pre-commitment. Facts are the residue of commitment, viewed ontologically. The same structure underlies all three.

The VERSF programme can therefore drop A1' as an axiom: it is now a theorem.

Acknowledgements

This work forms part of an ongoing programme on the foundational role of commitment events in physical theory. Earlier formulations of the commitment–entropy link relied on A1' as an axiom; the present derivation completes its reduction to a theorem.

References

- [1] Landauer, R. (1961). Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, **5**(3), 183–191.
- [2] Bennett, C. H. (1982). The thermodynamics of computation — a review. *International Journal of Theoretical Physics*, **21**(12), 905–940.
- [3] Bennett, C. H. (2003). Notes on Landauer's principle, reversible computation, and Maxwell's demon. *Studies in History and Philosophy of Modern Physics*, **34**(3), 501–510.
- [4] Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, **27**(3), 379–423; **27**(4), 623–656.
- [5] Jaynes, E. T. (1957). Information theory and statistical mechanics. *Physical Review*, **106**(4), 620–630.
- [6] Bekenstein, J. D. (1981). Universal upper bound on the entropy-to-energy ratio for bounded systems. *Physical Review D*, **23**(2), 287–298.
- [7] Wheeler, J. A. (1990). Information, physics, quantum: The search for links. In *Complexity, Entropy, and the Physics of Information* (W. H. Zurek, ed.), pp. 3–28. Westview Press.
- [8] Zurek, W. H. (2003). Decoherence, einselection, and the quantum origins of the classical. *Reviews of Modern Physics*, **75**(3), 715–775.

- [9] Joos, E., Zeh, H. D., Kiefer, C., Giulini, D., Kupsch, J., & Stamatescu, I.-O. (2003). *Decoherence and the Appearance of a Classical World in Quantum Theory* (2nd ed.). Springer.
- [10] Hardy, L. (2001). Quantum theory from five reasonable axioms. *arXiv preprint quant-ph/0101012*.
- [11] Chiribella, G., D'Ariano, G. M., & Perinotti, P. (2011). Informational derivation of quantum theory. *Physical Review A*, **84**(1), 012311.
- [12] Maroney, O. J. E. (2009). Generalizing Landauer's principle. *Physical Review E*, **79**(3), 031105.
- [13] Lloyd, S. (2002). Computational capacity of the universe. *Physical Review Letters*, **88**(23), 237901.
- [14] Maxwell, J. C. (1871). *Theory of Heat*. Longmans, Green, and Co.
- [15] von Neumann, J. (1955). *Mathematical Foundations of Quantum Mechanics* (R. T. Beyer, trans.). Princeton University Press. (Original work published 1932.)
- [16] Khinchin, A. I. (1957). *Mathematical Foundations of Information Theory* (R. A. Silverman & M. D. Friedman, trans.). Dover Publications.
- [17] Rényi, A. (1961). On measures of entropy and information. In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability*, Vol. 1, pp. 547–561. University of California Press.