The Information Substrate Theory
Volume VII

Information as the Name of Differentiatedness, and the Known Measures as Projections of One Substrate

Foundations, the six projections, and the bridge from loss to recovery to the boundaries of a subsystem's reach

Boris Kriger

Information Physics Institute
Institute of Integrative and Interdisciplinary Research, Toronto
ORCID: https://orcid.org/0009-0001-0034-2903

Seventh volume of the αLGQV / IST research programme

The Programme: The Seven Volumes

This volume is the seventh and final part of a single research programme. The first five volumes build a physics — from the cosmological scale inward to the strong sector — on one idea: that the structures we attribute to a dark sector, to singularities, and to confinement are consequences of the local gravitation of the quantum vacuum and of the topology of closed configurations. The sixth volume lifts that physics into a cross-scale method. The seventh, the present volume, supplies the lens: the information-theoretic ontology under which the whole is read.

The order is a descent followed by a turn. Volumes I–III move from the cosmos through extreme curvature to the empirical record; Volume IV provides the complex-algebraic apparatus the strong sector requires; Volume V carries the programme into quantum chromodynamics and the topology of nuclei on the physically correct manifold $\mathbb{RP}^3$; Volume VI generalises the topological apparatus of I–V into a universal cross-scale recognition method and an atlas of structure across fifteen scales. Volumes I–VI thus build and then survey an *Atlas* — a body of computed and measured structure from the nucleon to the cosmic web. Volume VII does not extend the Atlas; it reads it. It is placed last because it is a lens, not a foundation-stone: it interprets what the prior volumes established, showing each scale to be the materialization of a projection of one differentiated substrate. The programme is monolithic — not because every scale is the same, but because each realises a facet of one ground, and the seventh volume is the account of that ground and its facets.

Volume I — A Dark-Sector-Free Cosmology. *The Local Gravitation of Quantum Vacuum* (αLGQV). A vacuum–matter coupling derived from nuclear physics ($\alpha \approx 0.005$ predicted, 0.003 observed) reproduces flat rotation curves, the baryonic Tully–Fisher and radial-acceleration relations (with a_0 derived, not fitted), and dark-mass-free galaxies — from one parameter. doi:10.5281/zenodo.19027460

Volume II — Beyond Singularities. The programme at extreme curvature: toroidal vacuum cores replacing singularities, the interior of collapsed objects, the mass desert, the impossibility of wormholes, and the pole of spacetime replacing the Big Bang singularity, with a falsification programme for LIGO, LISA, JWST, EHT, and CMB-S4. doi:10.13140/RG.2.2.29913.28002

Volume III — Predictions Confirmed. The empirical compendium: NANOGrav, P-ACT, DESI dark-energy evolution, the Big Ring and Giant Arc, the cosmic radio dipole, CMB parity asymmetry, and galactic spin chirality — each independently predicted by αLGQV, each independently observed. doi:10.13140/RG.2.2.32586.32962

Volume IV — ICAST. *Imaginary Component Attribution Swaps Theory.* The complex-algebraic apparatus for σ-reattribution — the transfer of conserved quantum-number content between configurations during topological deformation — and the ontological backbone of the strong-sector programme. doi:10.5281/zenodo.20137134

Volume V — Topology of Quantum Chromodynamics and Isotopology. The strong sector: confinement as topological closure, mass as topology, nuclei as closed multibaryon configurations, magic numbers as symmetric topologies. A foundational paper argues the Yang–Mills mass gap is misposed on $\mathbb{R}^3$ and is a geometric consequence of compactness with antipodal identification on $\mathbb{RP}^3$. doi:10.5281/zenodo.20370383

Volume VI — Structural Persistence and Coherence Analysis (SPCA). The methodological extension: a cross-scale recognition framework built on a topological–dissipative analysis of stability and the Banach contraction to the unique fixed point x^*, with an atlas of ten worked recognitions from the nucleon to the global topology of the universe. Its discipline — honest acknowledgment of limits — is the one this volume inherits. doi:10.5281/zenodo.20455316

The present volume is **Volume VII — The Information Substrate Theory (IST)**: the lens that reads the Atlas of Volumes I–VI.

Synopsis of Volume VII

The foundational move. The volume begins from a claim it argues rather than assumes: that absolute nothing is impossible, so something exists, and that what exists is necessarily *differentiated* — an undifferentiated existent would be no existent at all. From this it draws its central definition: *information is the name of differentiatedness.* Information is not a substance the world is made of, not a fluid it contains, not a second ingredient alongside matter and energy; it is the word for the world's being differentiated. The axes along which we read the world — time, space, energy — and the laws stated in their terms are *facets* of the differentiated substrate, projections of it onto single coordinates. The substrate itself is *axis-free*: it carries no projection value, no number of bits, no entropy, no program length, because such values are readings along axes and the substrate is what is read, not a reading.

The one distinction. Everything in the volume turns on a single distinction — between the *substrate*, the differentiated and axis-free ground, and a *projection*, a reading of it along one axis — and on a single discipline: never mistake a projection for the substrate. This is the load-bearing line. Held, it yields the classification of the known information measures, the account of the model-building subsystem, and the two boundaries of knowledge. Dropped, the frame collapses into the reifications it exists to prevent: the world taken to *be* information, to *be* a computation, to *be* a simulation.

Part I — Foundations. The opening articles establish the substrate as the differentiated ground (information its name), distinguish substrate from projection and show the universe needs no information added to its physics, prove a *projection theorem* by which the known information measures descend from the substrate as readings along axes, and draw the *categorial boundary*: predicates internal to an axis (a colour, an age, a measure) do not apply to the axis-free substrate; questions that so apply them are not unanswered but ill-formed.

Part II — The Projections. The known theories of information are classified as projections along six axes, each a reduction that keeps one facet and discards the rest: the *channel* axis (Shannon), the *description* axis (Kolmogorov), the *regularity* axis (effective complexity), the *macrostate* axis (thermodynamic entropy), the *cost* axis (Landauer, and the infodynamics of DESI DR1), and the *receiver* axis (Bateson's difference — the one axis that carries no canonical number, the home of meaning). Inter-axis morphisms relate the six, and an empirical hinge — the Atlas article — shows each axis materialising at a definite scale of the prior volumes: the description axis as the nucleon and the atom, the regularity axis as knotted DNA and the living cell, the channel and macrostate axes as the cosmic web, and the global register as the $\mathbb{RP}^3$ topology of the universe. Throughout, the discipline holds: a projection is *realised* at a scale, never *identified* with the substrate.

Part III — The Bridge. The third part follows the consequence that organises the rest. Physical projection is many-to-one and therefore *loses* irreversibly — the data-processing inequality, the second law, and the uncomputability of inversion are its three faces — and the world cannot un-project. Yet structure projection discards comes to be known, because a model-building subsystem *recovers* it: not by inverting the projection, which is blocked, but by *inference*, maintaining a defeasible internal surrogate tested against the projections still to come. From this central asymmetry the bridge unfolds. The *arrow of time* is the subsystem's gradient of accumulated records; its direction comes not from the time-symmetric dynamics but from the low-entropy past, and the block does not flow — the "now" is the record frontier. The *simulation*, real as the limit of the recovering model, lives in the subsystem and not in the world; the question whether the world is a simulation, on the substrate reading, mislocates a projection. The *incompleteness of measurement* is a theorem: any proper projective access has a non-trivial joint fiber, so a subsystem necessarily has well-typed structure it can never resolve. And two boundaries bound a subsystem's reach — the *apophatic*, at the substrate, where questions cease to be well-formed, and the *projective*, within the world, where well-formed questions cease to be answerable — bounding, between them, a rich and knowable interior.

What the volume dissolves. A family of standing puzzles at the border of physics and philosophy is shown to be one confusion — a projection taken for the substrate — in several costumes, and dissolved together rather than one at a time: whether the universe computes (it need not; computation is a projection we apply), whether the world is made of information (no; information is the name of differentiatedness), whether the world is a simulation (the question is ill-typed), what the arrow of time is (the model-builder's record gradient), and why measurement cannot be complete (projective access has fibers).

What the volume is, and is not. The substrate frame is an *ontology* — a claim about the differentiated ground physics presupposes — and a *discipline* — a rule for speaking of it without reification. It is *conservative about physics*: it adds no equation, predicts no new number, and affirms physical theory entirely, changing the *status* of its quantities (projections, not substances) rather than their content. It is not a new physics, not a speculative metaphysics positing hidden entities, and neither idealism nor mysticism: the substrate's only ineffability is its axis-freeness, fully explicable. The volume's central open problem, posed and left open, is the *morphism geometry* — whether one diagram of the axes and their morphisms carries both the completeness of the classification and the incompleteness floor, as its image and its complement. And the frame is offered to be carried, reducible for use to one diagnostic question: does this predicate of the substrate something that belongs only to a projection?

How to Read This Volume

The volume is one argument, not a collection, and it is built to be read in order. Part I fixes the distinction on which all else turns; Part II is what the foundations make of the known measures; Part III is what the projections imply for a subsystem that has only projections to go on. Each part presupposes the one before, and a reader who accepts the foundations will find the projections and the bridge their consequences. Readers chiefly interested in the empirical anchoring may turn first to the Atlas article (10a), which binds the abstract axes to the computed structures of Volumes I–VI; readers interested in the synthesis may read Articles 16 and 17 last, where the whole is drawn together and placed. Each article is self-contained in statement, marks explicitly what it proves and what it conjectures, and ends with an epistemic-status ledger; the reader is never asked to take an entailment on trust.

The Articles

Part I — Foundations

Part II — The Projections

A classification of the known theories of information.

Part III — The Bridge

The Information Substrate:
Differentiation as the Ground of Information

Information as the name of differentiatedness, and the known measures as projections of it along single axes

Boris Kriger[1,2]

[1] Information Physics Institute
boris.kriger@informationphysicsinstitute.net

[2] Institute of Integrative and Interdisciplinary Research, Toronto
boriskriger@interdisciplinary-institute.org

ORCID: https://orcid.org/0009-0001-0034-2903

The Information Substrate Theory (IST) — Volume VII, Article 1
Part I: The Substrate — foundations of the volume

Abstract

The measures that go by the name *information* — Shannon's entropy, Kolmogorov complexity, the thermodynamic entropies, effective complexity, Landauer's bound, Bateson's difference — disagree about almost everything, yet each is called a measure of information. This opening article of the Information Substrate Theory asks what they are measures *of*, and answers that they are projections of a single feature of being: that it is differentiated. We record, from prior work in this programme and not re-argued here, that nothing is impossible; from this it follows that being cannot be undifferentiated, since the undifferentiated is indistinguishable from nothing. *Information* is then not a substance, a force, a constituent, or an act, but the *name* of differentiatedness — ineliminable from being not because the world requires it but because its absence is impossible. We give this claim its disciplined form. A differentiation structure carries only bare distinction; an *axis* is a reduction retaining the distinctions registered along one coordinate; the known measures are the values of canonical functionals on those reductions, while the substrate itself, carrying no axis, is measured by none. The claim is modal, not temporal: no act, agent, event, or moment appears, and because time is itself one axis, temporal predicates are categorially inapplicable to the substrate, a point we state and prove rather than assert. We locate the predecessors generously — each saw a facet — and we mark the boundary apophatically: within the describable the universe is a completed, invariant structure, and beyond it not even categories extend. An epistemic-status ledger separates what is established, what is here defined, and what later articles must earn, with the load-bearing derivation of the measures-as-projections claim assigned to the Projection Theorem of Article 3.

Keywords: information; differentiation; substrate; projection; foundations of physics; modality; Shannon entropy; Kolmogorov complexity; philosophy of information; apophatic boundary.

Contents

1 Introduction

There is no single quantity called information. There is a family of measures, each precise, each useful, and each at odds with the others about what information is. Shannon's entropy is a property of a probability distribution and grows with disorder [12]. Kolmogorov complexity is a property of an individual object, the length of its shortest description, and also grows with disorder, yet is computed by wholly different means and is uncomputable [5]. Effective complexity is *low* for both the perfectly ordered and the perfectly random and high only between them [2]. Thermodynamic entropy carries units of energy over temperature where the others are dimensionless. Landauer's bound attaches an energetic price to erasing a bit, importing a constant the combinatorial measures do not have [11]. Bateson's "difference that makes a difference" is not a number at all but a relation to a receiver [1]. The natural and usual conclusion is that "information" is an umbrella word for several unrelated tools.

This article proposes the contrary. The measures are not unrelated; they are projections of one structure seen along different axes. The structure is differentiatedness as such — the bare fact that being is differentiated rather than featureless — and each measure is what one obtains by reducing that structure to a single coordinate (a channel, a description, a macrostate, a regularity, a thermodynamic cost, a receiver) and applying the functional native to that coordinate. On this reading the disagreements among the measures are the differences among the *axes*, not among the things measured, and the word *information* names not any of the measures but their common source.

The aim of this opening article is to establish that source and to fix the vocabulary the rest of the volume uses. We do four things. We carry out the descent to the substrate: from the impossibility of nothing to the necessary differentiatedness of being to the identification of information as the name of that differentiatedness (Section 3). We give differentiatedness a minimal formal model and show that the apparatus of axes and laws — including time — consists of *facets* of the substrate rather than products of it (Section 4). We draw the central distinction of the whole theory, between the substrate and its projections, and state precisely the sense in which the substrate is measured by none of the known measures (Section 5). And we locate the predecessors, each of whom saw a facet truly (Section 6), before marking the boundary beyond which the theory is silent (Section 7).

A word on register, owed to the reader at the outset. This volume does not hold that matter is information, that information is fundamental, that the universe is a computer, or that the present proposal belongs in the lineage that reads physics as information-theoretic at bottom. It holds the smaller and, we think, more defensible claim that *information* is the name of an ineliminable feature of being, that every quantitative notion of information is a projection requiring a prior choice of axis, and that the substrate itself adds nothing to the world and is required by it in no form. The discipline of saying less and establishing it, rather than more and declaring it, is the discipline of the series, and it is nowhere more important than in its foundational article. The one genuinely load-bearing claim — that the known measures *descend* from the substrate by reduction along an axis — is not argued here but in the Projection Theorem of Article 3 of this volume, which this article sets up and on which it relies.

2 Logical Structure of the Argument

The argument is a descent followed by a distinction, then a placement and a boundary. The reader may hold the following map in view.

Step 1 (Section 3): The descent.
From the impossibility of nothing — taken as established in prior work, not re-argued — we obtain that being is necessarily differentiated, because the undifferentiated is indistinguishable from nothing (Proposition 3.1). We then fix the meaning of *information* as the name of differentiatedness: not a substance, act, or residue, but the word for being's ineliminable feature.

Step 2 (Section 4): The structure and its facets.
We model differentiatedness as a differentiation structure carrying only bare distinction, and argue that axes (time, space, energy) and laws (regularities) are facets of that structure, not products of it in time. Because time is itself one axis, temporal predicates are categorially inapplicable to the substrate (Proposition 4.5); the descent is modal, not temporal.

Step 3 (Section 5): Substrate and projection.
We define an axis as a reduction retaining one coordinate's distinctions, and a projection as the substrate taken along an axis. The known measures are projections; the substrate, carrying no axis, is in the domain of no axis-bound measure (Proposition 5.2). That the named measures actually descend from the substrate in this way is the theorem of Article 3, to which we defer the derivation.

Step 4 (Section 6): Locating the predecessors.
Each major conception of information — Shannon, Wheeler, Landauer, Bateson, Spencer-Brown, Jaynes, Gell-Mann and Lloyd — is placed as the disclosure of a facet, correct within its axis. The stance is to locate, never to demolish.

Step 5 (Section 7): The boundary.
Within the describable the universe is a completed, invariant structure; beyond it the theory is silent, and silent not only of entities but of categories. The boundary of the describable is the boundary of predicate applicability.

Step 6 (Section 8): Epistemic status.
An explicit ledger separates what is established, what is here defined, and what later articles must earn.

The dependencies are linear: Step 2 presupposes Step 1; Step 3 presupposes Step 2; Steps 4 and 5 presuppose Step 3; Step 6 audits the whole. The thesis advanced by the totality is that information is the name of differentiatedness, and the plurality of its measures the plurality of axes along which differentiatedness can be taken.

3 The Descent to the Substrate

3.1 There is something

We begin from a result we do not re-argue in the main descent. Prior work in this programme has defended the impossibility of nothing: the claim that the null state of affairs — the absence of any being whatever — is not a possible state of affairs but a pseudo-possibility that dissolves on inspection [6, 8]. We take this as the established premise of the present descent and build from it. Because the entire volume rests on this premise, however, we do not leave it merely a pointer: a self-contained statement of what the premise means and of the argument for it — reconstructed from the prior work and marked, as a philosophical argument should be, with its own premises — is given in Appendix B. The reader who wishes to see exactly what must be granted is referred there; the main descent uses only the conclusion.

> **Premise (established elsewhere).** There is something rather than nothing, and necessarily so: nothing is impossible.

We note, against a common misreading, that the premise is not the empirical observation that the world happens to be non-empty. It is the modal claim that emptiness is not among the alternatives — that "nothing" does not name a rival state the world might have been in. This distinction matters for everything that follows, because it forbids us from treating being as a contingent addition to a prior void. There was no void to which being was added; the void is not a was.

3.2 Being is necessarily differentiated

From the impossibility of nothing we derive the first positive feature of being. We require a minimal notion of distinction, which we will formalize in Section 4 but state informally now: being is *undifferentiated* if it sustains no distinction — if there is, within it, no telling of anything from anything.

Proposition 3.1 (The undifferentiated is indistinguishable from nothing)**.** *Let a candidate state of being sustain no distinction whatever. Then it is indistinguishable from the null state. Consequently, since the null state is impossible, being cannot be undifferentiated.*

Proof. Suppose a state sustains no distinction. Then for any putative feature, predicate, or part, there is no telling its presence from its absence or one of its values from another, for to tell them apart would be to sustain a distinction, contrary to hypothesis. A state in which no presence is distinguishable from any absence is one in which nothing can be said to obtain rather than not obtain; it answers every question of the form "is this so?" with no determinate difference. But this is precisely the condition of the null state, which likewise sustains no obtaining of anything. Hence the undifferentiated and the null are indistinguishable: there is no distinction in virtue of which they would differ. Since, by the established premise, the null state is impossible, any state indistinguishable from it is impossible. Therefore being, which is not impossible, is not undifferentiated; it is differentiated. □

The conclusion is exactly as strong as its premise and no stronger. We have not shown that being *must contain* this or that structure, only that it cannot lack distinction altogether, on pain of collapsing into the impossible null. The content of the differentiation — which distinctions, how organized — is not fixed by the argument and is not the substrate's to fix in advance; it is what the axes, introduced later, render visible. What the argument fixes is solely the *ineliminability* of differentiatedness: being is differentiated not as a property it contingently has but as the minimal condition of its being anything at all — differentiatedness as the ontological condition of actualization, in the sense developed in prior work of the programme [7].

3.3 Information is the name of differentiatedness

We can now say what *information* means in this theory, and, just as importantly, what it does not.

Definition 3.2 (Information, in the substrate sense)**.** *Information* is the name of differentiatedness: the term by which we refer to the fact that being is differentiated rather than featureless. It is not a substance, a field, a force, a conserved quantity, a fluid, an act, an event, or a residue left by a process. It adds nothing to being and is required by being in no form; it is the word for a feature being cannot lack.

The definition is deliberately deflationary, and the deflation is the point. To say that being is "informational" on this account is not to add information to the inventory of the world's contents, alongside matter and energy; it is to observe that whatever the world's contents are, they are differentiated, and *information* is our name for their being so. The universe does not run on information, store information, or compute; these are predicates of certain *projections* (Section 5) and of certain subsystems that build models, not of the substrate. The substrate neither possesses information as a stock nor processes it as a flow. It is differentiated, and that is all that "information," at the level of the substrate, says.

This is the place to be explicit about the tension a reader of the series will feel. Earlier volumes argued that *information* is a spectator-relative, dimensionless, and multiply defined word that ought not to be admitted as a fundamental physical concept, recommending *constraint* in its stead [10]. The present volume is titled for information.

There is no reversal. Definition 3.2 does not promote information to a fundamental constituent; it identifies information with differentiatedness, which is not a constituent at all but a feature, and it then shows — this is the work of Section 5 and of Article 3 — that every *quantitative* notion of information is a projection requiring a prior choice of axis. The spectator-relativity the earlier volumes flagged is real and is here given its location: it enters with the choice of axis, not with the substrate. The substrate is spectator-free precisely because it is axis-free.

3.4 Counterarguments and Responses

Objection 1 (the descent smuggles in an observer through "distinction"). To speak of distinction, the objection runs, is to presuppose someone who distinguishes; the argument therefore covertly assumes a subject and is not about being as such. *Response.* Distinction in Proposition 3.1 is not an act of distinguishing but a standing feature: that there *is* something to be told apart, whether or not anyone tells it. The proof uses only the counterfactual "were one to ask, there would be a determinate difference," which is a fact about the state, not about an asker. An observer is required to *register* a distinction along an axis (Section 5), and that is where spectator-relativity correctly enters; it is not required for the distinction to obtain.

Objection 2 (the impossibility of nothing is contestable, so the whole descent is fragile). *Response.* It is contestable, and we have not defended it here; we have imported it. The honesty demanded by the series requires us to mark the descent as conditional on that premise, which we do in the ledger of Section 8. What we claim is the conditional: *if* nothing is impossible, *then* being is differentiated and the rest follows. A reader who rejects the premise rejects the descent, and is owed the prior work, not a repetition of it.

Objection 3 (calling differentiatedness "information" is a mere relabeling that explains nothing). *Response.* By itself the name would indeed explain nothing; a name is not a theory. The explanatory content lies entirely downstream, in the claim that the known measures are projections of the so-named substrate (Article 3) and that the relations among them are the rigidities of the morphisms between projections. The name earns its keep only if that downstream work succeeds. We make the name's modesty explicit here so that it is not mistaken for the theory; the theory is the descent of the measures, not the christening of the substrate.

3.5 Section Result and Implications

We have established, conditional on the imported impossibility of nothing, that being is necessarily differentiated (Proposition 3.1), and we have fixed *information* as the name of that differentiatedness, expressly denying it the status of substance, act, or fundamental constituent (Definition 3.2). The implication carried forward is that the question "what is information?" has, at the level of the substrate, a one-word answer —

differentiatedness — and that all further content lies in how differentiatedness is *taken*: along which axis, by which reduction, under which functional. To that structure we now turn.

4 The Differentiation Structure and Its Facets

4.1 The minimal model

Within the describable, we model the substrate by the least structure that distinguishes being from nothing: a carrier of bare distinction and nothing more.

Definition 4.1 (Differentiation structure)**.** A *differentiation structure* is a pair $\mathcal{S} = (\Omega, \sim)$ where Ω is a nonempty set of configurations and $\sim$ is the identity relation, regarded as the finest distinction: $x \sim y$ iff $x = y$. We require $|\Omega| \geq 2$. The structure carries no metric, no measure, no ordering, and no coordinate; it records only that distinct configurations are distinct.

The requirement $|\Omega| \geq 2$ is the formal residue of Proposition 3.1: a one-element carrier sustains no distinction and is, by that proposition, the formal counterpart of nothing. We record this as a corollary, since it licenses the lower bound.

Proposition 4.2 (Triviality of the undifferentiated model)**.** *If* $|\Omega| = 1$*, then every reduction of* $\mathcal{S}$ *(Definition 4.3) is constant, and* $\mathcal{S}$ *is indistinguishable under all axes from the empty structure.*

Proof. With a single configuration there is no pair $x \neq y$ to distinguish, so any map retaining distinctions retains none and is constant; no axis registers anything. Indistinguishability from the empty structure follows as in Proposition 3.1. □

We stress that Definition 4.1 is a model *within the describable*. It is not a claim that the substrate "is" a set, with all a set's further properties; it is the thinnest mathematical object adequate to carry the one feature the descent established, namely bare distinguishability. Everything beyond distinguishability that a set ordinarily brings — cardinality arithmetic, a particular topology — is to be regarded as not yet present at the level of the substrate and as entering only with an axis.

4.2 Axes and laws as facets

Definition 4.3 (Axis and reduction)**.** An *axis* α on $\mathcal{S}$ is a surjection $\rho_\alpha \colon \mathcal{S} \to \mathcal{S}_\alpha$ onto an axis-reduced structure, which retains exactly the distinctions registered by the coordinate α and collapses all others. The image $\mathcal{S}_\alpha$ is the α-*projection* of the substrate.

With this we can state the structural claim that distinguishes IST from the manifesto picture it declines to join. In the manifesto picture, time, space, and energy are the stage on which differentiation plays out, and laws are the rules it obeys; differentiation is then a process *in* the axes. In the substrate picture the order is reversed: the axes and the laws are *facets* of differentiatedness — coordinates along which the one differentiated structure can be taken — and not a prior arena within which it occurs.

Remark 4.4 (Facets, not products)**.** To call time, space, energy, and the regularities *facets* is to say that each is a way the substrate is differentiated, recoverable as a reduction ρ_α, rather than a thing produced by the substrate at some stage. The claim is structural and carries no temporal index: there is no moment before which the axes were absent and after which they appeared, because "before" and "after" are themselves the time-facet and have no application to the substrate as a whole.

This brings us to the most delicate point of the article, which we state as a proposition rather than leave as an attitude.

Proposition 4.5 (Categorial inapplicability of temporal predicates to the substrate)**.** *Let time be the axis α_t with reduction $\rho_{\alpha_t} \colon \mathcal{S} \to \mathcal{S}_{\alpha_t}$. Then temporal predicates — "before," "after," "when," "how long," "once or repeatedly" — are defined only on the time-projection $\mathcal{S}_{\alpha_t}$ and not on $\mathcal{S}$. To apply them to the substrate itself is a category error: the substrate is not in the domain of these predicates.*

Proof. A temporal predicate is, by construction, a function of the time-coordinate; its arguments are points or intervals of $\mathcal{S}_{\alpha_t}$. The substrate $\mathcal{S}$ is not such a point or interval but the un-reduced structure of which $\mathcal{S}_{\alpha_t}$ is one projection under ρ_{α_t}. Applying a function to an argument outside its domain is undefined; the application of a temporal predicate to $\mathcal{S}$ is therefore not false but ill-typed. The same holds, *mutatis mutandis*, for any axis-specific predicate applied to the substrate rather than to its corresponding projection. □

The proposition is the rigorous form of the discipline announced at the outset. It forbids us from saying that being "becomes" differentiated, that differentiation "happens," that it occurs "once" or "continuously," or that there was a "first" differentiation — not because these claims are false but because they are ill-formed, applying the time-facet to that of which time is only a facet. The descent of Section 3 is to be read throughout in this light: "being is differentiated" is modal and structural, with no hidden verb of process.

4.3 Counterarguments and Responses

Objection 1 (modeling the substrate as a set already imports too much). *Response.* The model imports only what is used: a carrier and the identity relation, with the explicit instruction (after Definition 4.1) to treat all further set-theoretic structure as axis-introduced rather than intrinsic. Where a later argument uses more than bare distinguishability — a probability, a length, a count — it does so only after an axis has supplied it, and the supply is marked. The set is scaffolding for the describable model, not a metaphysical claim that being is a set.

Objection 2 (if time is only a facet, the theory cannot accommodate real change). *Response.* It accommodates change exactly where change lives: on the time-projection $\mathcal{S}_{\alpha_t}$, where temporal predicates are defined and where the ordinary physics of process proceeds untouched. What the theory forbids is not change within the time-facet but the projection of change onto the substrate as a whole, which would be to ask when the whole came to be differentiated — a question Proposition 4.5 shows to be ill-typed. The block of the describable contains all its change; it does not itself change, because it is not in the time it contains.

Objection 3 (calling laws "facets" dissolves their explanatory force). *Response.* It relocates, rather than dissolves, that force. A law, as a regularity, is a recurring distinction — a facet retained under the regularity axis (Article 3 treats this axis as the home of effective complexity). To call it a facet is to say it is a way the substrate is differentiated, which is precisely what makes it available to explain: regularities explain because they recur, and they recur because they are standing facets, not because they are imposed from outside on a prior arena.

4.4 Section Result and Implications

We have given differentiatedness a minimal model (Definition 4.1), shown the undifferentiated model to be trivial (Proposition 4.2), defined an axis as a reduction (Definition 4.3), and established that axes and laws are facets of the substrate rather than products of it, with the sharp consequence that temporal and other axis-specific predicates are categorially inapplicable to the substrate itself (Proposition 4.5). The implication is that the substrate is the axis-free common origin of every coordinate, and that the apparatus of physics — time, space, energy, law — stands to it as projection stands to source. The next section makes this substrate–projection relation the central distinction of the theory.

5 Substrate and Projection

5.1 The central distinction

The whole theory turns on a single distinction, which we can now state cleanly.

Definition 5.1 (Substrate and projection)**.** The *substrate* is differentiatedness as such, the structure $\mathcal{S}$ prior to any axis. A *projection* is the substrate taken along an axis, i.e. an image $\rho_\alpha(\mathcal{S}) = \mathcal{S}_\alpha$. A *projection measure* is a functional M_α on $\mathcal{S}_\alpha$; the corresponding *information measure* is $M_\alpha \circ \rho_\alpha$.

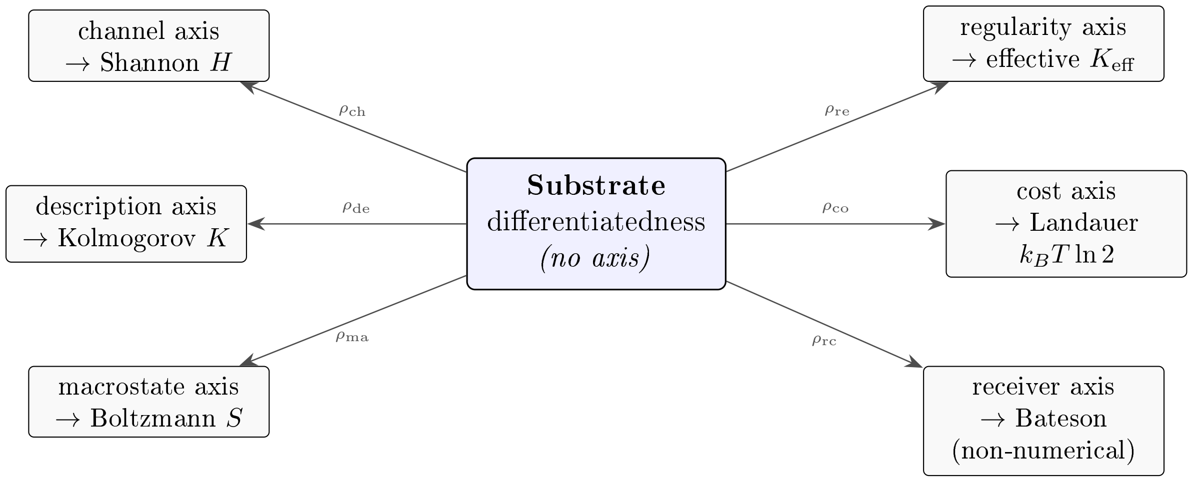

Figure 1: The central image of the volume. One substrate — differentiatedness as such, carrying no axis — and the known information measures as its projections along single axes. The reductions ρ_α retain the distinctions registered by one coordinate and collapse the rest; the substrate itself is in the domain of no projection measure. Article 3 proves that the named measures descend in exactly this way, deriving Shannon and Kolmogorov from their reductions and locating the remainder.

The known measures populate this scheme as projection measures along distinct axes. We list them, with the axis each occupies, deferring to Article 3 the demonstration that they genuinely descend in this way; here the list fixes vocabulary and orients the reader.

- **Channel axis** — the ensemble of distinguished outcomes with weights; projection measure: Shannon entropy.
- **Description axis** — a single configuration as a finite object to be reconstructed; projection measure: Kolmogorov complexity.
- **Macrostate axis** — a coarse-graining and its microstate count; projection measure: Boltzmann–Gibbs entropy.
- **Regularity axis** — the recurring facet of a configuration; projection measure: effective complexity.
- **Cost axis** — the minimum thermodynamic price of a logical operation; “measure”: Landauer’s bound (with an external dimensional bridge).
- **Receiver axis** — the distinctions a given receiver registers; “measure”: Bateson’s difference (not numerical).

5.2 The substrate is measured by none of its projection measures

The distinction has an immediate and important consequence, which is the foundational counterpart of the locality clause of the Projection Theorem.

Proposition 5.2 (The substrate carries no projection value)**.** *For every axis α, the projection measure M_α is a function of $\mathcal{S}_\alpha = \rho_\alpha(\mathcal{S})$, not of $\mathcal{S}$. The substrate $\mathcal{S}$ is therefore in the domain of no M_α, and no number that any known information measure assigns is a number assigned to the substrate.*

Proof. By Definition 5.1, M_α takes arguments in $\mathcal{S}_\alpha$. The substrate $\mathcal{S}$ is not an element of $\mathcal{S}_\alpha$ but the domain of the surjection ρ_α whose codomain is $\mathcal{S}_\alpha$; it is the source, not a point of the image. Hence $M_\alpha(\mathcal{S})$ is undefined for every α, and any value $M_\alpha(\rho_\alpha(\mathcal{S}))$ is a value of the projection $\mathcal{S}_\alpha$, carried by the axis, not by $\mathcal{S}$. $\square$

This is why the question "how much information does the universe contain?" has, on this theory, no answer at the level of the substrate, only answers relative to a chosen axis — a number of bits on the channel axis, a description length on the description axis, an entropy on the macrostate axis — none of which is the substrate's own. The substrate is not a large reservoir of information; it is differentiated, and "how much information" is a question one can ask only after choosing how to take it.

5.3 The asymmetric bridge, in preview

One feature of the substrate–projection relation must be flagged here, though its development belongs to Part III of the volume. Projection is *lossy*: a reduction collapses many substrate configurations to one projected value, and the collapse is not in general reversible. A subsystem that models the world must therefore *recover* structure from impoverished projections, reconstructing by hypothesis what the projection discarded. This asymmetry — physical projection loses, the model recovers — is, in the later articles, the source of the arrow of time as a projection phenomenon, of the incompleteness of measurement, and of the sense in which a sufficiently complex model becomes a simulation that lives in the subsystem and not in the world. We record the asymmetry now only to mark that the substrate–projection relation is directional, and to forestall the misreading that projection is a mere change of coordinates with nothing lost. It is not; what it loses is the subject of the bridge articles.

5.4 Counterarguments and Responses

Objection 1 (the substrate, being measured by nothing, is idle — an unobservable posit doing no work). *Response.* The substrate does precisely one kind of work, and it is not measurement: it is the common domain whose existence makes the relations among the projection measures consequences rather than coincidences. That work is discharged in Article 3, where the agreement of expected description length with entropy, the dissociation of effective complexity from Kolmogorov complexity, and the proportionality of thermodynamic to channel entropy are shown to be the rigidities of the morphisms between projections of one domain. An unobservable that forces observable relations is not idle; it is explanatory in the only way a common source can be.

Objection 2 (positing a single substrate behind many measures is unfalsifiable metaphysics). *Response.* The substrate *as named* is metaphysics, and we do not pretend otherwise; what is falsifiable is the structural claim it underwrites — that every robust relation among the measures is induced by a morphism between their reductions (Article 3). A robust cross-axis relation that provably cannot be so induced would refute the structural claim and with it the usefulness of the posit. The posit earns its place by the falsifiable work it does, not by its own observability.

Objection 3 (the axis list is arbitrary; why these six?). *Response.* The list is not claimed complete or canonical, and its open-endedness is marked as conjecture in Article 3's ledger and in ours. Each axis is specified by an independent reduction with its own characterization in the prior literature, not reverse-engineered to host a measure; that several measures sort cleanly onto independent axes is evidence for the scheme, not proof of its completeness, and we claim only the former.

5.5 Section Result and Implications

We have drawn the central distinction between the substrate and its projections (Definition 5.1), placed the known measures as projection measures along distinct axes (Figure 1), and proved that the substrate is in the domain of no projection measure (Proposition 5.2). We flagged the lossy, directional character of projection as the seed of the bridge articles. The implication, carried to the rest of the volume, is that "information" splits cleanly into two registers: the substrate sense, in which it is the axis-free name of differentiatedness, and the projection sense, in which it is one of several axis-bound measures. Conflating the two registers is the characteristic error the theory is built to avoid, and Article 9 will display a worked instance of it.

6 Locating the Predecessors

The classification stance of this volume is that the major conceptions of information are each correct within an axis, and that the task is to locate them rather than to adjudicate among them. The diversity that looked like confusion in Section 1 is, on the substrate reading, a catalogue of facets truly seen.

Shannon [12] saw the channel axis with unmatched clarity, attaching the measure of registered difference to a probability distribution over distinguishable messages. The often-noted observer-dependence of Shannon information is, in the present vocabulary, the dependence of the channel *reduction* on a prior choice of which outcomes count as distinct; the functional adds no observer, and Shannon's theory is exactly right about the channel-facet.

Spencer-Brown [13] took the act of drawing a distinction as primitive and built a calculus on it. This is the substrate's own primitive — bare distinguishability — approached from the side of the operation rather than the standing feature, and it is the nearest predecessor to Definition 4.1.

Bateson [1] located information on the receiver axis: a difference that makes a difference to a system organized to register it. The conception is correct and is why the receiver axis appears in our scheme; its non-numerical character is not a defect but the honest mark of an axis that carries no canonical functional.

Jaynes [3, 4] read thermodynamic entropy as an inference from incomplete distinction among microstates, which is precisely the macrostate-axis reading; the maximum-entropy principle is the statement that, absent further facets, the macrostate reduction selects the least committal microstate weights.

Landauer [11] disclosed the cost axis, binding a logical operation on distinctions to a minimum thermodynamic price through an external constant $k_B T$. The insight is genuine and is located, in our scheme, as a projection requiring a dimensional bridge rather than as a property of the substrate.

Gell-Mann and Lloyd [2] isolated the regularity axis, distinguishing the description of an object's recurring structure from the description of its random part, and thereby gave effective complexity its characteristic non-monotonicity — which Article 3 reads as the signature of a refinement morphism.

It remains to be generous, and exact, about the conception this volume most often is mistaken for. Wheeler [14] proposed that physics is, at bottom, information-theoretic — "it from bit" — and the proposal has inspired a large and serious literature. On the substrate reading, Wheeler saw something true: that distinction is not incidental to physics but constitutive of it. Where the present volume parts company is in register, not in respect. "It from bit" makes the bit — a projection-axis quantity — fundamental, and reads being off the channel facet; IST declines the inversion, holding that the bit is a projection and that what is fundamental, if anything is, is the axis-free differentiatedness of which the bit is one reduction. The difference is exactly the difference between promoting a facet to the whole and naming the whole of which it is a facet. We record Wheeler's program with respect and locate it as the channel facet over-read as ontology — a characteristic and instructive over-reading, not an error of observation.

6.1 Counterarguments and Responses

Objection 1 (locating everyone as "right within an axis" is uncritical ecumenism). *Response.* The ecumenism is bounded and does real work. To locate a conception on an axis is a falsifiable placement, not a compliment: it predicts which relations that conception can and cannot capture, and it is refuted if the conception captures cross-axis relations that its axis cannot reach. The stance is charitable about *observation* — each predecessor saw a facet — and exacting about *scope*, which is where the locating bites.

Objection 2 (the charitable reading of Wheeler is a way of claiming his insight while disowning his thesis). *Response.* We claim no insight of Wheeler's as our own; we credit the insight to him and dispute only the inversion that makes the bit fundamental. The distinction is principled: one may hold that distinction is constitutive of physics (with Wheeler) while denying that any particular projection-axis quantity is the ground (against "it from bit"). That is not appropriation but disagreement about register, stated as such.

6.2 Section Result and Implications

We have located the principal conceptions of information — Shannon, Spencer-Brown, Bateson, Jaynes, Landauer, Gell-Mann and Lloyd, and Wheeler — each as the disclosure of a facet, correct within its axis and instructive in its scope. The implication is that the substrate reading is not a rival to these conceptions but a frame that holds them, assigning to each its coordinate and reserving for the substrate only what no facet captures: the axis-free differentiatedness they each project. This is the sense in which IST aims to be a theory *of* the measures rather than another measure among them.

7 The Boundary

A foundational article owes a statement of where its theory stops, and the statement must be as disciplined as the descent. We give it here in brief; Article 4 of the volume is devoted to it.

Within the describable, the universe on this account is a completed, invariant structure — a block whose differentiatedness is given, whose axes are its facets, and whose change lives on the time-facet without the whole being in time (Proposition 4.5). The block is realist: its differentiatedness does not depend on being projected, and the projection measures report facts about its facets, not artifacts of the reporting.

Beyond the describable the theory is silent, and silent in a strict sense that bears emphasis. The silence is not only of *entities* — we do not furnish the beyond with a simulator, a creator, a multiverse, or a first cause — but of *categories*. The boundary of the describable is the boundary of predicate applicability: temporal predicates, causal predicates, and the very predicate "exists" are facets or projections and have no purchase beyond the structure of which they are facets. To ask what is "outside" the block, or what was "before" it, or what "caused" it, is to apply axis-bound predicates beyond their domain, and is ill-typed for the same reason that asking when the substrate became differentiated is ill-typed. The theory does not answer these questions; it declines them as malformed, which is a different and more disciplined thing than leaving them open.

This apophatic boundary is not a concession but a load-bearing commitment. It is what prevents the substrate from being quietly refurnished as a metaphysical entity behind the world — a god, a computer, a ground of being in the theological sense. The substrate is the axis-free differentiatedness of the describable block, no more; it is not a thing the block sits on, and there is, in the strict sense above, no place for it to sit.

7.1 Section Result and Implications

We have marked the boundary: within the describable, a completed, invariant, realist block; beyond it, the silence not only of entities but of categories, with the boundary of the describable identified as the boundary of predicate applicability. The implication is that the theory's modesty is structural, not merely temperamental: it cannot overreach into the beyond without committing the same category error it forbids at the substrate, and its discipline is therefore enforced by its own form. Article 4 develops this boundary in full; the present statement fixes it for the foundational record.

8 Epistemic Status: Established, Defined, To Be Earned

In the manner of the series, we close with an explicit ledger. For a foundational article the relevant statuses are slightly different from those of a theorem-bearing one. *Established (imported)* marks premises taken from prior work and not re-argued here. *Established (here)* marks the small results actually proved in this article. *Defined* marks vocabulary fixed by definition, which is sound but carries no empirical claim. *To be earned* marks claims this article relies on but defers to a named later article, and *Conjectured* marks claims the programme advances but has not established anywhere.

Claim	**Status**	**Locus**
Nothing is impossible	Established (imported); restated	Kriger [8]; App. B
Being is necessarily differentiated	Established (here)	Prop. 3.1
Undifferentiated model is trivial	Established (here)	Prop. 4.2
Temporal predicates inapplicable to substrate	Established (here)	Prop. 4.5
Substrate carries no projection value	Established (here)	Prop. 5.2
Information = name of differentiatedness	Defined	Def. 3.2
Differentiation structure; axis; projection	Defined	Defs. 4.1, 4.3, 5.1
Known measures descend from the substrate	To be earned	Article 3 (Projection Theorem)

continued on next page

Epistemic-status ledger (continued)

Claim	Status	Locus
Relations among measures are induced	To be earned	Article 3
Projection is lossy; the model recovers	To be earned	Part III (bridge articles)
Boundary is apophatic (categories do not cross)	To be earned	Article 4
Mistaking a projection for the substrate (worked case)	To be earned	Article 9
The axis family is complete / canonical	Conjectured	Article 3 ledger

The ledger makes the division of labor explicit. This article *proves* only what bare distinction and the imported premise yield: the necessity of differentiatedness and the categorial facts about the substrate. It *defines* the vocabulary. And it is candid that its central promise — that the known measures descend from the substrate, and that their relations are thereby induced — is not discharged here but in Article 3, which this article exists to set up.

9 Discussion

The result of this opening article is a vocabulary and a descent, not a theorem about the measures; that theorem is Article 3's. What the article secures is the frame within which the rest of the volume is intelligible: that there is a substrate, that it is differentiatedness as such, that *information* names it without adding to it, and that the familiar measures are its projections along axes the substrate does not itself carry. The frame is built to be modest in three specific ways, each of which is a deliberate refusal of the manifesto register. It refuses to make information a constituent of the world, holding it to be the name of a feature. It refuses to make any projection-axis quantity — a bit, an entropy — fundamental, locating each as a reduction. And it refuses to furnish the beyond, holding the boundary of the describable to be the boundary of predicate applicability.

The frame also resolves, rather than incurs, the tension that a reader of the series will have felt at the title. The earlier volumes' refusal to admit information as a fundamental physical concept is not reversed but vindicated: the substrate sense of information is not a physical concept at all but the name of a feature, and the projection sense is exactly the spectator-relative, axis-bound, dimensionless thing the earlier volumes flagged. To name the volume for information is therefore not to elevate information but to study, with discipline, what the word can and cannot bear.

9.1 Limitations

Three limitations are material and are marked rather than hidden. *First*, the descent is conditional on an imported premise, the impossibility of nothing; a reader who rejects that premise rejects the descent, and this article does not re-argue it. *Second*, the central explanatory claim of the frame — that the known measures descend from the substrate and that their relations are induced — is asserted here but proved only in Article 3; read alone, this article fixes vocabulary and proves only the categorial facts about the substrate. *Third*, the substrate as *named* is a metaphysical posit, and its warrant is entirely the falsifiable structural work it underwrites downstream; an article that stopped at the naming would have established a frame and no physics.

9.2 Future Directions

The immediate sequel is internal to the volume and already largely in hand. Article 2 argues in detail that the universe needs no information in any form, sharpening Definition 3.2 against the storage-and-flow picture. Article 3 discharges the load-bearing promise, deriving Shannon and Kolmogorov as projections and proving that the relations among the measures are the rigidities of the morphisms between reductions. Article 4 develops the apophatic boundary of Section 7 into a full treatment of the predicates that do not cross the block. Beyond the volume's own architecture, the open question this article bequeaths is the one its scheme cannot yet settle: whether the axes admit an intrinsic classification — a principle generating the admissible reductions of a differentiation structure — which would turn the conjectured completeness of the axis family into a theorem or refute it. That question is stated here and pursued where the apparatus to address it exists.

Acknowledgments

The author thanks colleagues of the Information Physics Institute and the Institute of Integrative and Interdisciplinary Research for sustained discussion of the foundations of the series.

References

[1] G. Bateson. *Steps to an Ecology of Mind*. Chandler Publishing, San Francisco, 1972.

[2] M. Gell-Mann and S. Lloyd. Information measures, effective complexity, and total information. *Complexity*, 2(1):44–52, 1996.

[3] E. T. Jaynes. Information theory and statistical mechanics. *Physical Review*, 106(4):620–630, 1957.

[4] E. T. Jaynes. Information theory and statistical mechanics. II. *Physical Review*, 108(2):171–190, 1957.

[5] A. N. Kolmogorov. Three approaches to the quantitative definition of information. *Problems of Information Transmission*, 1(1):1–7, 1965.

[6] B. Kriger. The law of limit to negation: negation can't negate itself. Institute of Integrative and Interdisciplinary Research. Zenodo, 2025. doi:10.5281/zenodo.18101306.

[7] B. Kriger. Differentiation as the ontological condition of actualization. Zenodo, 2026. doi:10.5281/zenodo.18268520.

[8] B. Kriger. Testing nothingness and evaluating the inevitability of life: a deductive chain from being to persistence, intelligence, and the limits of inquiry. Institute of Integrative and Interdisciplinary Research. Zenodo, 2026. doi:10.5281/zenodo.18764754.

[9] B. Kriger. The categorial boundary: on predicates that do not cross the block. The Information Substrate Theory (IST), Volume VII, Article 4. Zenodo, 2026.

[10] B. Kriger. *Structural Persistence and Coherence Analysis: A Topological–Dissipative Method Across Physical Scales*. Monograph series, Volume VI. Zenodo, 2026. doi:10.5281/zenodo.20455316.

[11] R. Landauer. Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5(3):183–191, 1961.

[12] C. E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423, 1948.

[13] G. Spencer-Brown. *Laws of Form*. Allen & Unwin, London, 1969.

[14] J. A. Wheeler. Information, physics, quantum: the search for links. In W. H. Zurek, editor, *Complexity, Entropy, and the Physics of Information*, pages 3–28. Addison-Wesley, 1990.

A The Descent as a Numbered Chain

For the record we restate the descent of Sections 3–5 as a single numbered chain of premises and conclusions, with the status of each step marked. This is the skeleton on which the prose of the article hangs; it is offered so that a reader may check the dependencies without reconstructing them.

(D1) *(Premise, imported; restated in App. B.)* Nothing is impossible: the null state of affairs is not a possible state of affairs [6, 8].

(D2) *(Definition.)* A state is undifferentiated if it sustains no distinction.

(D3) *(Lemma, here.)* An undifferentiated state is indistinguishable from the null state (proof of Proposition 3.1).

(D4) *(Conclusion, here.)* Being is not undifferentiated; being is necessarily differentiated (Proposition 3.1).

(D5) *(Definition.)* *Information* is the name of differentiatedness, and is neither substance, act, nor constituent (Definition 3.2).

(D6) *(Definition.)* An axis is a reduction retaining one coordinate's distinctions; a projection is the substrate taken along an axis (Definitions 4.3, 5.1).

(D7) *(Conclusion, here.)* Time is one axis; hence temporal predicates are categorially inapplicable to the substrate (Proposition 4.5). The descent is modal, not temporal.

(D8) *(Conclusion, here.)* The substrate is in the domain of no projection measure (Proposition 5.2).

(D9) *(Deferred to Article 3.)* The known information measures descend from the substrate by reduction along an axis, and the relations among them are induced by the morphisms between reductions.

(D10) *(Deferred to Article 4.)* The boundary of the describable is the boundary of predicate applicability; beyond it the theory is silent of categories.

Steps (D1)–(D8) are settled within this article (D1 by import and the restatement of Appendix B, the rest by the proofs cited). Steps (D9)–(D10) are the load-bearing deferrals on which the article's explanatory value depends and which the named later articles discharge.

B What "Nothing Is Impossible" Means: A Self-Contained Statement

The descent of Section 3 imports, from prior work, that nothing is impossible [6, 8]. Because the entire volume rests on this premise, we set down here, in one place, what the premise means and what the argument for it is. The statement is philosophical, not mathematical, and we mark its own premises as such, so that a reader can see exactly what must be granted and exactly where a critic would press. We do not claim the argument is beyond dispute; we claim it is valid given the premises stated, that those premises are highly plausible, and that this is the claim on which the volume builds.

B.1 What "nothing" purports to name

The word "nothing" has an innocent, restricted use and a contested, absolute one, and only the absolute use is at issue.

Definition B.1 (Restricted and absolute nothing)**.** *Restricted nothing* is the absence of some kind within a wider being: nothing in the box, nothing in the visible spectrum, no even prime above two. Restricted nothing is unproblematic — it is a fact about a region of being — and the impossibility claim does not concern it. *Absolute nothing* purports to name the total absence of being: no entities, no properties, no relations, no facts, and no truths — not even the truth that nothing obtains. The impossibility claim concerns absolute nothing alone.

The distinction matters because objections to "nothing is impossible" almost always trade, illicitly, on the restricted sense, where the claim is obviously false (there is plenty of restricted nothing). The thesis is about the absolute sense, and its content is that the absolute sense fails to name a possible state of affairs.

B.2 The self-undermining of absolute negation

The core of the argument is the law of limit to negation [6]: a negation cannot negate itself; an absolute, unrestricted negation runs out before it reaches its own obtaining. We reconstruct it as a reductio, with the two premises it requires made explicit.

Proposition B.2 (Impossibility of absolute nothing)**.** *Assume:*

(N1) (Obtaining entails the fact of obtaining.) *If a state of affairs obtains, then it is the case that it obtains — there is the fact, or truth, that it obtains.*

(N2) (A fact is a truth-bearer.) *That a state of affairs is the case is itself something true: where a state obtains, at least one truth holds, namely that it does.*

Then absolute nothing, as defined in Definition B.1, does not obtain in any possible state of affairs; it is not a possible state.

Proof. Suppose, for reductio, that absolute nothing obtains. By (N1), it is then the case that absolute nothing obtains; by (N2), this is a truth — the truth that absolute nothing obtains. Hence, on the supposition, there is at least one truth. But by Definition B.1, absolute nothing is the total absence of truths, including the truth that nothing obtains; so on the supposition there is no truth. The supposition therefore yields both that there is a truth and that there is none, a contradiction. Hence absolute nothing does not obtain in any possible state of affairs; it is not a possible state. This is the precise sense of the law of limit to negation: the absolute negation of being would have to negate even the truth of its own obtaining, which it cannot do while obtaining; the negation cannot reach itself. □

The conclusion is modal and exactly as strong as (N1)–(N2). It does not say that there is an impossible object called "nothing"; it says that "absolute nothing" fails to name a possible state, that the void is not among the alternatives. Equivalently: necessarily, there is something. There is no possible world that is empty in the absolute sense, because the supposed emptiness would have to include its own truth and exclude it at once.

B.3 What the argument does and does not establish

Three boundaries keep the result from being read as more than it is.

First, it establishes the impossibility of the absolute void, not the necessity of any particular being. That there must be something does not entail that there must be *this* something, or that anything in particular exists necessarily; the modal force is de dicto (necessarily, something exists) and not de re (of some thing, it necessarily exists). The argument fixes only that being has no empty alternative, which is all the descent of Section 3 uses.

Second, it offers no cause or explanation of being. To ask what *makes* it the case that there is something, or what the absence of the void is *due to*, is to seek a ground behind being; Article 4 shows such causal extrapolation past the describable to be ill-typed [9]. The argument removes the void as an alternative; it does not install a reason in its place.

Third, and most important for honesty, the argument rests on (N1) and (N2), and a determined critic will press exactly there. One might deny (N2) — deny that "nothing obtains" is a genuine truth-bearer, holding instead that in the absolute void there is nothing to bear truth and hence no truth that nothing obtains, so that no contradiction arises. The reply, from [6], is that this denial does not describe a coherent state but withdraws from description altogether: a putative state of which it cannot even be said that it obtains is not a state that has been specified as obtaining, and the friend of the absolute void wanted a state that obtains. The dialectic does not terminate in a knock-down, and we do not pretend it does; we record that the volume's premise is this argument, that its load is carried by (N1)–(N2), and that a reader who rejects (N1)–(N2) rejects the descent and is owed this appendix rather than a louder assertion.

B.4 Relation to the undifferentiated and to the categorial boundary

Two connections place the result within the volume.

The impossibility of absolute nothing re-grounds the impossibility of the undifferentiated. Proposition 3.1 showed that a state sustaining no distinction is indistinguishable from the null state; Proposition B.2 shows the null state to be impossible; together they yield that the undifferentiated is impossible, which is the hinge of the descent. The present appendix thus supplies, in the volume itself, the premise that Proposition 3.1 discharges.

Finally, the impossibility of absolute nothing must not be confused with the categorial boundary of Article 4 [9], though both keep being from resting on a void. Absolute nothing is a would-be *global state* — the whole of being empty — and is shown *impossible*, a modal verdict. The "beyond the block" of Article 4 is a would-be *region or locative* — something outside the differentiated whole — and is shown *non-denoting*, a categorial verdict: not impossible but ill-typed. The two verdicts are different and are reached by different means, and the volume needs both: the first removes the empty alternative to being, the second removes the external region in which a ground might be lodged. Conflating them produces the ill-typed sentence "there is nothing beyond the block," which Article 4 forbids; holding them apart is the discipline that lets the volume say, without contradiction, both that there could not have been nothing and that "what is beyond" does not denote.

Substrate and Projection:
Why the Universe Needs No Information

Storage, processing, conservation, and reification as predicates of projections, not of the differentiated whole

Boris Kriger[1,2]

[1] Information Physics Institute
boris.kriger@informationphysicsinstitute.net

[2] Institute of Integrative and Interdisciplinary Research, Toronto
boriskriger@interdisciplinary-institute.org

ORCID: https://orcid.org/0009-0001-0034-2903

The Information Substrate Theory (IST) — Volume VII, Article 2
Part I: The Substrate — foundations of the volume

Abstract

It is widely said that the universe stores information, processes it, conserves it, or is in some sense made of it. This article argues that the universe needs information in none of these senses, and that each picture, properly located, is a statement about a *projection* rather than about the differentiated whole. The argument rests on the foundation of the preceding articles: information is the name of differentiatedness, and differentiatedness is ineliminable from being because its absence — the undifferentiated — is indistinguishable from nothing, which is impossible. We first make precise what "needs" means, and show that to need something is for its absence to be a possible state of failure; since the absence of differentiatedness is not a possible state at all, the universe does not *need* differentiatedness but cannot *lack* it, a stronger and quite different relation. We then take the four standard pictures in turn. The storage picture treats information as a stock, but a stock is the value of a projection measure — a count on the channel or macrostate axis — and the substrate carries no such value. The processing picture treats the universe as a computer, but processing is a time-indexed operation, and temporal predicates are categorially inapplicable to the substrate; computation is a predicate of model-building subsystems, not of the whole. The conservation picture — unitarity, and the black-hole information puzzle — is a true statement about a projection's content under time evolution, not about the substrate, which is not in time; we relocate the puzzle as a tension between two projections without presuming to resolve the physics. The reification pictures, from "it from bit" to mass–energy–information equivalence, promote a projection-axis quantity to the ground, which is the characteristic category error the theory is built to detect. We locate each predecessor generously, classify the pictures by the error each commits, and mark the boundary of what is here established.

Keywords: information; substrate; projection; storage; computation; conservation; unitarity; Landauer's principle; holographic bound; foundations of physics.

Contents

1 Introduction

A familiar way of speaking has the universe in the information business. It is said to *store* information — so many bits within a horizon, on a black hole's area, in the configuration of its matter. It is said to *process* information — to compute its own next state, a vast cellular automaton or quantum circuit grinding forward. It is said to *conserve* information — unitary evolution losing nothing, the black-hole information puzzle owing its sting to the conviction that information cannot truly be destroyed. And, at the most ambitious, it is said to *be* information — "it from bit," the world made of yes-or-no answers, or information endowed with mass and energy of its own. These ways of speaking are not idle; they have produced real physics and real mathematics, and the present article does not dispute that physics. It disputes the picture the speaking encourages: that information is something the universe possesses, manipulates, keeps, or is composed of — something it *needs*, in some form, to be what it is.

The thesis of this article is that the universe needs information in no form. The claim is not that information is unimportant or that the measures are wrong; it is that the relation between the universe and information is not the relation of a system to a resource it requires. The foundation laid in the preceding articles makes this precise. Information is the name of differentiatedness [5]; the universe is differentiated not because it requires differentiation to function but because the alternative, the undifferentiated, is indistinguishable from nothing, and nothing is impossible. Differentiatedness is therefore ineliminable from being, which is a stronger and quite different relation than being needed. A thing that is needed could, in some possible state, be absent, with failure as the consequence; differentiatedness has no such possible absence. The universe does not run on information the way an engine runs on fuel. It is differentiated, and "information" is our name for its being so.

Once this is seen, the four standard pictures sort themselves. Each, we will argue, is true of a *projection* — the substrate taken along a single axis — and becomes false, or rather ill-formed, only when transferred from the projection to the substrate. A stock of bits is a value of a projection measure, and the substrate carries no projection value (Section 4). Processing is a time-indexed operation, and the substrate is not in time (Section 5). Conservation is the constancy of a projection's content across time, a true statement about a projection and not about the timeless whole (Section 6). And the reification of information — making a bit, or an erasure cost, the ground of the world — promotes a projection-axis quantity to the substrate, which is the characteristic category error the whole theory is built to detect (Section 7). We classify the pictures, locate their distinguished authors generously, and are careful to mark where we relocate a problem without claiming to solve it.

A note on what this article does and does not do. It does not resolve the black-hole information paradox, and it says so plainly; it relocates the paradox, identifying it as a tension between two projections rather than a fact about the substrate. It does not refute the holographic bound or Landauer's principle; it endorses them as true statements about projections and denies only their transfer to the substrate. And it does not pretend that "the universe needs no information" is a discovery of new physics; it is a clarification of what the existing physics is about, which is the proper work of a foundational article.

2 Logical Structure of the Argument

The argument is one modal result followed by four diagnoses and a classification.

Step 1 (Section 3): What "needs" means.
To need something is for its absence to be a possible state of failure. Since the absence of differentiatedness is impossible (the foundation of Article 1), the universe does not need differentiatedness; it cannot lack it (Proposition 3.2). This fixes the sense of the title.

Step 2 (Section 4): The storage picture.
A stock of information is a value of a projection measure; the substrate carries no projection value (Proposition 4.1). The universe stores no bits.

Step 3 (Section 5): The processing picture.
Processing is a time-indexed operation; temporal predicates are categorially inapplicable to the substrate (Proposition 5.1). Computation is a predicate of model-building subsystems, not of the whole.

Step 4 (Section 6): The conservation picture.
Conservation is the constancy of a projection's content across time; it is true of a projection and inapplicable to the substrate (Proposition 6.1). The black-hole information puzzle is relocated as a tension between two projections, not resolved.

Step 5 (Section 7): The reification pictures.
Making a bit, an erasure cost, or an information-mass the ground of the world promotes a projection-axis quantity to the substrate, the characteristic category error. Each reification is located by its axis; the holographic and Bekenstein bounds are endorsed as projection bounds. A classification table summarizes the diagnoses.

Step 6 (Section 8): Epistemic status.
An explicit ledger separates what is established here, what is imported, and what is deferred.

The dependencies are linear: Steps 2–5 each instantiate, for one picture, the substrate–projection distinction fixed in Article 1 and the modal result of Step 1. The thesis advanced by the totality is that the universe's relation to information is not need, possession, or composition, but the ineliminability of a feature whose quantitative measures are all projections.

3 What "Needs" Means

3.1 Need and ineliminability

The title of this article makes a claim about *need*, and the claim is exact only if "need" is. We give it the natural modal analysis.

Definition 3.1 (Need)**.** A system X *needs* Y if there is a possible state of X in which Y is absent and in which X thereby fails to be, or to function as, X. Need thus presupposes that the absence of Y is a genuine possibility for X.

Proposition 3.2 (The universe does not need differentiatedness; it cannot lack it)**.** *If the absence of Y is not a possible state of X, then X does not need Y in the sense of Definition 3.1; rather Y is* ineliminable *from X. The absence of differentiatedness is the undifferentiated, which is indistinguishable from nothing and hence impossible [5]. Therefore the universe does not need differentiatedness — equivalently, does not need information in the substrate sense — but cannot lack it.*

Proof. By Definition 3.1, need requires a possible Y-absent state of failure. If no such state is possible, the antecedent of need is unsatisfiable, and the need relation does not hold; what holds instead is ineliminability, the impossibility of the Y-absent state. By the foundation of the preceding article, the undifferentiated is impossible, so the differentiatedness-absent state of the universe is impossible. Hence the universe does not need differentiatedness and cannot lack it. Since information, in the substrate sense, is the name of differentiatedness, the same holds of information. □

The distinction between need and ineliminability is not pedantry; it is the whole point. To say the universe *needs* information is to picture a universe that might, in some circumstance, run short of it, as an engine runs short of fuel or a computer of memory — and then to ask how the supply is maintained, where the bits come from, what would happen if they ran out. Proposition 3.2 removes the picture at its root. There is no circumstance of shortage, because there is no possible state of absence; the question "what supplies the universe's information?" is as ill-posed as "what supplies its being-something-rather-than-nothing." The universe is differentiated in the way it is something: not by provision but by the impossibility of the alternative.

3.2 Why the rest of the article is necessary

Proposition 3.2 settles the substrate sense of information at one stroke, but it does not by itself dispose of the pictures of Section 1, because those pictures rarely speak of the substrate sense. They speak of *quantities*: a number of bits stored, a rate of operations, a conserved measure, a fundamental unit. These are projection-sense notions, and to show that the universe needs none of them in any form, we must show, picture by picture, that each is a predicate of a projection rather than of the substrate. That is the work of the next four sections. The modal result fixes the substrate sense; the diagnoses fix the projection senses; together they discharge the title.

3.3 Counterarguments and Responses

Objection 1 (this is a verbal trick: "ineliminable" just is "needed in the strongest way"). *Response.* The two are not the same and the difference is operational. "Needed" invites the questions of supply, shortage, and maintenance — where does it come from, can it run out — and these questions have answers, or at least sense, for things that are needed. "Ineliminable" forecloses those questions as ill-posed: there is no supply because there is no possible absence to be supplied against. A relation that changes which questions are well-formed is not a mere relabeling of another relation.

Objection 2 (Proposition 3.2 leans entirely on the impossibility of nothing, which is contestable). *Response.* It does, and we mark it as imported in the ledger of Section 8, exactly as Article 1 did. The conditional we assert is robust whatever one thinks of the premise: *if* the undifferentiated is impossible, *then* the universe does not need differentiatedness. A reader who rejects the premise owes a different account of why there is something rather than nothing, but the structure of the present argument is unaffected.

3.4 Section Result and Implications

We have fixed the sense of "need" (Definition 3.1) and shown that the universe does not need differentiatedness but cannot lack it (Proposition 3.2), so that the substrate sense of information is settled: the universe needs it in no form because need has no purchase where the alternative is impossible. The implication is that any remaining force in the pictures of Section 1 must come from their projection-sense content — from quantities — which we now examine one picture at a time.

4 The Storage Picture

4.1 Stocks are projection values

The storage picture holds that the universe contains an amount of information — that there is a number of bits which is the information content of a region, a horizon, or the whole. The number is often large and sometimes precise. The claim we make is not that the number is wrong but that it is a number belonging to a projection.

Proposition 4.1 (A stock of information is a projection value)**.** *Any "amount of information" — a count of bits, a number of distinguishable states, an entropy — is the value $M_\alpha(\rho_\alpha(\mathcal{S}))$ of a projection measure along some axis α (channel, description, or* *macrostate). By the locality of projection measures, established in Article 1, the substrate $\mathcal{S}$ is in the domain of no M_α. Hence no stock of information is a property of the substrate; the substrate stores no amount.*

Proof. Each candidate notion of an amount is, by inspection, a functional on an axis-reduced structure: a count of distinguishable outcomes is a functional on the channel projection, a microstate count is a functional on the macrostate projection, a description length is a functional on the description projection. By the locality result of the foundational article [5], these functionals take arguments in $\mathcal{S}_\alpha$, not in $\mathcal{S}$; $M_\alpha(\mathcal{S})$ is undefined. Therefore any stock is a value carried by a projection, not by the substrate. □

The point is sharpened by an example the storage picture itself supplies. The Bekenstein bound [1] and the holographic principle [10] state that the entropy — hence, on the macrostate axis, the number of distinguishable states — of a bounded region is limited by its surface area in Planck units. This is a true and deep statement, and on the present reading it is a statement about a projection: it bounds the value of the macrostate-axis projection measure for a region under gravitational constraints. It does not say that the substrate of the region "holds" that many bits as a stock; it says that the macrostate projection of the region cannot exceed a certain value. Read as a fact about a projection, the bound is exactly right and loses none of its force. Read as a fact about the substrate — as though the region were a container filled to a maximum with bit-stuff — it commits the locality error of Proposition 4.1.

4.2 Counterarguments and Responses

Objection 1 (the holographic bound is manifestly about physical content, not a mere bookkeeping projection). *Response.* It is about physical content, and the present reading affirms this: the macrostate projection of a physical region is a physical fact about that region, not bookkeeping. What the reading denies is only the further step, that the bounded quantity is a stock the substrate *contains* as it might contain a fluid. The bound constrains a projection; it does not stock a substrate. The physics is untouched; only the picture of bit-stuff in a container is removed.

Objection 2 (if the substrate has no amount, IST cannot even agree that one region is "more differentiated" than another). *Response.* It can, but only via a projection. "More differentiated" is itself a comparative that requires an axis to be made quantitative; along the macrostate axis one region may project to a larger count than another. The comparison is real and is a comparison of projections. What has no axis-free answer is the absolute question "how much information does the region contain, full stop," which is precisely the question Proposition 4.1 shows to be ill-posed at the level of the substrate.

4.3 Section Result and Implications

We have shown that any stock of information is the value of a projection measure and hence not a property of the substrate (Proposition 4.1), and we have located the Bekenstein and holographic bounds as true statements about the macrostate projection rather than about substrate content. The implication is that the universe "stores" information only in the sense that its projections take values; it holds no reservoir of bits. We turn from the static picture of a stock to the dynamic picture of a process.

5 The Processing Picture

5.1 Processing is a temporal predicate

The processing picture holds that the universe computes — that it carries out, step by step, operations on its information, as a cellular automaton or a quantum circuit does. The picture has been made quantitative: one may estimate the number of elementary operations the universe could have performed in its history [9], and the estimate is a legitimate fact about a physical system evolving in time. The claim we make is that this fact, too, belongs to a projection, and specifically to the time-facet.

Proposition 5.1 (Processing is inapplicable to the substrate)**.** *"Computation," "transmission," "processing," and cognate predicates are time-indexed: their arguments are states at times and transitions between them, i.e. elements and ordered pairs of the time-projection $\mathcal{S}_{\alpha_t}$. By the categorial inapplicability of temporal predicates to the substrate, established in Article 1, these predicates do not apply to $\mathcal{S}$. The substrate computes, transmits, and processes nothing; such predicates are well-formed only on the time-projection and of the subsystems situated there.*

Proof. A processing predicate asserts that one state is transformed into another across a step, which presupposes a temporal ordering of states; its arguments therefore lie in the time-projection. The substrate is not in the time-projection but is the unreduced structure of which the time-projection is one facet [5]; applying a time-indexed predicate to it is ill-typed, as for any axis-specific predicate applied to the substrate. Hence processing predicates apply to subsystems on the time-projection, not to the substrate. □

There is a constructive counterpart to this negative result, which the bridge articles of Part III develop and which we state here only to locate it. Projection is lossy: the time-facet, like every reduction, discards structure. A subsystem that must act within the time-facet on impoverished projections is driven to *recover* the discarded structure by hypothesis — to build a model that predicts, and a sufficiently elaborate predictive model is what we recognize as computation. On this reading the computation in the world is real but located: it lives in the model-building subsystems, which are in time and which process projections, not in the substrate, which is not in time and processes nothing. The universe is not a computer; it contains computers, which are subsystems recovering what projection lost. The estimate of the universe's "computational capacity" [9] is then read, without subtraction, as a bound on the processing that time-bound subsystems could have carried out, which is a fact about the time-projection and the physics within it.

5.2 Counterarguments and Responses

Objection 1 (physical law evolves states in time, so the universe just does compute, by definition). *Response.* Physical law evolves states on the time-projection, and there the language of computation is apt; the present reading agrees that the time-facet supports dynamics and, with model-building subsystems, computation. What it denies is the extension of "computes" from the time-projection to the substrate as a whole, which would be to ask what the whole computes across a time it is not in. The dynamics is real where time is real — in the projection — and the substrate is not an additional, larger computer behind it.

Objection 2 (calling computation a "subsystem" predicate is anthropocentric; nature computes without us). *Response.* No observer is required by the reading. A model-building subsystem need not be a mind; it is any physical subsystem whose dynamics, on the time-projection, reconstructs structure that projection discarded, as the bridge articles make precise. The point is not that computation needs us but that it needs a *subsystem in time*, which the substrate is not. Nature computes wherever there are such subsystems; it does not compute as a whole, because the whole is not in time.

5.3 Section Result and Implications

We have shown that processing predicates are time-indexed and hence inapplicable to the substrate (Proposition 5.1), and we have located real computation in the model-building subsystems that inhabit the time-projection. The universe is not a computer; it contains them. The implication carried forward is that both the static picture (storage) and the dynamic picture (processing) attribute to the substrate what belongs to its projections. The subtlest of the pictures, conservation, attributes to the substrate what is in fact a constancy of a projection across time, and we treat it with corresponding care.

6 The Conservation Picture

6.1 Conservation is the constancy of a projection across time

The conservation picture holds that information cannot be destroyed — that under the exact dynamics of the world, the information of a closed system is preserved. In quantum theory this is the unitarity of evolution, and it is the conviction behind the sense of paradox in the fate of information that falls into a black hole and appears to be lost in thermal Hawking radiation [4]. The physics here is serious and we do not contest it. We contest only the location of the conserved quantity.

Proposition 6.1 (Conservation is a time-facet predicate)**.** *"Conserved" means "constant across the time-coordinate." Its subject is therefore a quantity defined on the time-projection and its predicate is constancy along α_t. The statement "information is conserved" is thus a true statement, where it holds, about the constancy of a* projection's content *under time evolution (for instance, the channel content under unitary dynamics). It is not, and cannot be, a statement about the substrate, to which the temporal predicate "constant across time" does not apply (Article 1).*

Proof. Constancy across time is a predicate whose evaluation requires comparison of a quantity at different times, hence arguments in the time-projection. Unitary evolution preserves inner products and so the channel/description content of a closed system along the time-coordinate; this is a fact about the time-projection of that system. The substrate is not in the time-projection [5]; "the substrate conserves information" applies a temporal predicate to a non-temporal subject and is ill-typed. Hence conservation, where true, is a statement about a projection's constancy, not about the substrate. □

6.2 The black-hole information puzzle, relocated

The black-hole information puzzle is, on the standard framing, the apparent conflict between two well-supported claims: that the radiation from an evaporating black hole is thermal and carries no information about what formed it [4], and that the underlying evolution is unitary and so preserves it. We do not resolve this conflict, and we are explicit that the present framework does not entitle anyone to a resolution. What the framework offers is a relocation.

Remark 6.2 (The puzzle as a tension between two projections)**.** On the substrate reading, the two claims concern two different projections of the same underlying differentiated process: a macrostate-axis projection (the thermal spectrum and the entropy of the radiation) and a channel/unitary-axis projection (the preserved content under exact evolution). The puzzle is the question whether these two projections are jointly consistent — whether the joint structure of the two reductions, in the sense of Article 3, admits both the thermality of the one and the preservation of the other. So framed, the puzzle is neither a fact about the substrate gaining or losing information (the substrate, being timeless, neither gains nor loses) nor a paradox of the kind that a single projection could exhibit; it is a question about the compatibility of two projections of one process. Whether they are compatible — and by what physics — is exactly the open problem, and the framework neither prejudges it nor claims to settle it.

The value of the relocation is diagnostic, not dispositive. It says that the *form* of the puzzle is a two-projection consistency question, which clarifies what a resolution must do: it must exhibit the joint structure of the macrostate and unitary projections and show how both predicates can hold of it. It does not say which proposed resolution is correct, and it would be a violation of the discipline of this volume to suggest that renaming the puzzle dissolves it. The puzzle is real physics; the framework only locates the registers in which its two horns live.

6.3 Counterarguments and Responses

Objection 1 (if conservation is "only" about a projection, IST is denying unitarity). *Response.* It denies nothing of the sort. Unitarity is affirmed as a true and important fact about the time-projection of closed quantum systems; the conserved content is a projection's content, and its conservation is a theorem of the dynamics on the time-facet. The only claim added is that this conservation is not a property of the timeless substrate, which is not the kind of thing that is constant or inconstant across a time it is not in. Affirming unitarity and denying that the substrate is its subject are consistent.

Objection 2 (relocating the black-hole puzzle is evasion dressed as analysis). *Response.* It would be evasion if it were offered as a solution, and we are at pains to say it is not. A relocation earns its place only by clarifying the form of the problem, and the clarification here is specific and falsifiable in spirit: it predicts that any genuine resolution will take the shape of a joint-structure consistency between the macrostate and unitary projections, and it would be embarrassed by a resolution that fit no such description. That is a claim about the form of admissible answers, not a refusal to engage.

6.4 Section Result and Implications

We have shown that conservation is the constancy of a projection's content across time and is inapplicable to the timeless substrate (Proposition 6.1), and we have relocated the black-hole information puzzle as a tension between two projections of one process, without presuming to resolve it (Remark 6.2). The implication is that even the most physical-seeming sense in which the universe "keeps" information is a sense in which a projection keeps it across time; the substrate keeps nothing, because keeping is temporal. We turn last to the pictures that make information not a stock, a process, or a conserved content, but the very ground of the world.

7 The Reification Pictures

7.1 Promoting a projection to the ground

The reification pictures are the most ambitious and, in the present diagnosis, commit the most direct version of a single error. They take a projection-axis quantity and make it fundamental: a bit, in "it from bit"; an erasure cost, in the strongest readings of Landauer's principle; an information-mass, in proposed mass–energy–information equivalences. The error is uniform and we state it as such.

Proposition 7.1 (The reification error)**.** *Let $q = M_\alpha(\rho_\alpha(\mathcal{S}))$ be the value of a projection measure along an axis α. To assert that q, or the unit in which q is counted, is a fundamental constituent or property of the substrate is to attribute to $\mathcal{S}$ a quantity that, by the locality of projection measures (Article 1), is not defined on $\mathcal{S}$. The assertion is therefore not false but ill-typed: it predicates of the substrate a value belonging to one of its projections.*

This is the same locality error met in Proposition 4.1, now committed not about a stock but about the ground. We take the three principal reifications in turn, locating each generously.

It from bit. Wheeler [13] proposed that physics is at bottom information-theoretic, that every physical entity derives its existence from apparatus-elicited yes-or-no answers — bits. The insight, as Article 1 already granted, is that distinction is constitutive of physics, not incidental to it. The reification is the further step that makes the bit — a value on the channel axis — the ground. On the present reading, distinction *is* fundamental, in the sense that being is differentiated; but the bit is not distinction, it is the channel-axis projection of distinction, and to ground the world in the bit is to ground it in one facet rather than in the differentiatedness of which the bit is a facet. We keep Wheeler's insight and decline his inversion.

The computational universe. The view that the universe is a computer [9] reifies the processing picture: it makes computation, a time-facet predicate, the nature of the whole. Section 5 located real computation in time-bound subsystems; the reification extends it to the substrate, committing Proposition 5.1's error in the ground register rather than the predicate register. The computational capacity of the universe is a fine fact about the time-projection; the universe's *being* a computer is the projection promoted to ground.

Mass–energy–information equivalence. The proposal that information carries mass or energy of its own, and the associated "second law of infodynamics" [11, 12], reifies the cost axis. Landauer's principle [2, 8], now experimentally confirmed in a single-bit memory [3], binds the erasure of a bit to a minimum dissipated energy $k_B T \ln 2$ through an external dimensional bridge; it is a true and measured fact about the cost axis. The reification gives the bit itself a mass, transferring the cost-axis quantity to the substrate as a substance. We do not assess the empirical claims of the mass–energy–information program here; that assessment, against established physics, is the dedicated subject of Article 9 of this volume. We diagnose only the category: a cost-axis projection, requiring a dimensional bridge to acquire energetic meaning at all, is being read as a property of the substrate, which is Proposition 7.1's error in its starkest form.

7.2 The five pictures and their diagnoses

The diagnoses of this article may be collected. Each picture is true of a projection and ill-typed of the substrate; the table records, for each, the projection it really concerns and the category error it commits when transferred to the substrate.

Picture	**Real locus (projection)**	**Diagnosis**
Storage ("N bits within a horizon")	channel / macrostate count	A stock is a projection value; the substrate carries none (Prop. 4.1). Bekenstein/holographic bounds are true of the projection.
Processing ("the universe computes")	time-projection; subsystem models	Processing is time-indexed and inapplicable to the timeless substrate (Prop. 5.1); real computation lives in subsystems.
Conservation ("information is never lost")	channel content under unitary time	Conservation is constancy across time, true of a projection, not of the substrate (Prop. 6.1); black-hole puzzle relocated, not resolved.
Reification ("it from bit"; computational universe)	channel axis; time-projection	A projection-axis quantity promoted to the ground (Prop. 7.1); insight kept, inversion declined.
Reification (mass–energy–information)	cost axis + dimensional bridge	A cost-axis projection given substance; category diagnosed here, empirical assessment deferred to Article 9.

7.3 Counterarguments and Responses

Objection 1 (the reification diagnosis would forbid any claim that information is physically real). *Response.* It forbids no such claim at the level of projections, where information-quantities are as physically real as the projections that bear them: the erasure cost is real and measured [3], the holographic bound is real, the channel content is real. What the diagnosis forbids is only the transfer of these projection-quantities to the substrate as constituents or properties of it. Information is physically real exactly where it is a projection value; it is not an additional substance of which the substrate is composed.

Objection 2 (deferring the mass–energy–information assessment to Article 9

is a way of avoiding the hardest case). *Response.* The division is deliberate and is not avoidance. The present article makes a *categorial* claim — that the proposal reads a cost-axis projection as substrate substance — which is independent of whether its specific empirical predictions succeed or fail. Article 9 makes the *empirical* assessment against established physics, including the cosmological data. Keeping the categorial and empirical questions separate is a condition of treating each fairly; collapsing them here would prejudge the data, which we decline to do.

7.4 Section Result and Implications

We have shown that the reification pictures commit a single locality error — promoting a projection-axis quantity to the ground (Proposition 7.1) — and we have located "it from bit," the computational universe, and the mass–energy–information program each by its axis, keeping the genuine insight of each and declining only the transfer to the substrate. The classification table collects the five diagnoses. The implication, completing the article, is that across all four pictures — storage, processing, conservation, reification — the universe's relation to information is never that of a system to a needed, possessed, kept, or constitutive resource; it is, in every case, the relation of a differentiated whole to the values its projections take.

8 Epistemic Status: Established, Imported, Deferred

In the manner of the series, we close with an explicit ledger. *Established (here)* marks results proved in this article; *Established (imported)* marks premises and results taken from prior articles and not re-argued; *Affirmed (physics)* marks established physical facts the article endorses and only relocates; *Deferred* marks claims handled in a named later article; *Relocated, not resolved* marks a problem whose register the article clarifies without solving.

Claim	**Status**	**Locus**
Undifferentiated is impossible	Established (imported)	Kriger [5]
Locality of projection measures	Established (imported)	Kriger [5]
Temporal predicates inapplicable to substrate	Established (imported)	Kriger [5]
Universe does not need, cannot lack, differentiatedness	Established (here)	Prop. 3.2
A stock of information is a projection value	Established (here)	Prop. 4.1
Processing is inapplicable to the substrate	Established (here)	Prop. 5.1
Conservation is a time-facet predicate	Established (here)	Prop. 6.1
Reification is a uniform locality error	Established (here)	Prop. 7.1
Bekenstein / holographic bound	Affirmed (physics)	Bekenstein [1], Susskind [10]
Unitarity of closed quantum evolution	Affirmed (physics)	standard

continued on next page

Epistemic-status ledger (continued)

Claim	**Status**	**Locus**
Landauer bound (measured)	Affirmed (physics)	Bérut et al. [3], Landauer [8]
Black-hole information puzzle	Relocated, not resolved	Rem. 6.2
Mass–energy–information empirical assessment	Deferred	Article 9
Computation in model-building subsystems	Deferred (mechanism)	Part III

The ledger marks the article's discipline. It *proves* the modal result and the four locality diagnoses; it *affirms* the established physics it relocates, adding no dispute to unitarity, the holographic bound, or the measured Landauer cost; it *relocates* the black-hole puzzle without claiming to resolve it; and it *defers* the empirical assessment of the mass–energy–information program to the article equipped to make it.

9 Discussion

The result of this article is a clarification, not a discovery: the universe needs information in no form, because the substrate sense of information is ineliminable rather than needed (Proposition 3.2), and every projection sense — stock, process, conserved content, fundamental unit — is a predicate of a projection rather than of the substrate. The four diagnoses are four instances of the single distinction the volume is built on, applied to the four ways the information-bearing picture of the universe is usually expressed. What the diagnoses remove is a picture, not a physics: the picture of the universe as a thing that holds, manipulates, keeps, or is made of information, with all the attendant questions of supply, capacity, loss, and composition. What they leave standing is the physics in full — the holographic bound, the computational capacity estimate, unitarity, the measured Landauer cost — relocated to the projections where it lives.

The clarification also continues the series' resolution of its own apparent tension. To title a volume for information and then argue that the universe needs none of it is not a contradiction once the two senses are kept apart. In the substrate sense, information is the ineliminable name of differentiatedness, needed in no form because its absence is impossible. In the projection sense, information is the family of axis-bound measures, each real, each physical, none a property of the substrate. The earlier volumes' refusal to make information a fundamental physical concept is thereby reaffirmed at the level of the substrate and located precisely at the level of the projections, where the spectator-relative, dimensionless, multiply defined character of the word is exactly the mark of its being a projection of something axis-free.

9.1 Limitations

Three limitations are material. *First*, the modal result rests on the imported impossibility of nothing; a reader who rejects that premise rejects the strongest form of the thesis and is left only with the locality diagnoses, which stand on their own but deliver "the universe possesses no axis-free amount" rather than "needs none." *Second*, the relocation of the black-hole information puzzle is diagnostic only; it clarifies the form a resolution must take and contributes nothing to the physics of the resolution itself, and it would be a misreading to take the relocation as progress on the paradox. *Third*, the treatment of the mass–energy–information program is deliberately confined to its category; the empirical question, which is the one most likely to decide the program's fate, is deferred to Article 9 and is not prejudged here.

9.2 Future Directions

The article's sequels are internal to the volume. Article 3 supplies the theorem the locality diagnoses presuppose, deriving the projection measures from their axes and proving that the relations among them are induced by the morphisms between reductions; the present diagnoses are, in effect, applications of that theorem's locality clause to the four pictures. Article 4 develops the apophatic boundary that the timeless, need-free reading of the substrate requires. Article 9 makes the empirical assessment of the infodynamics program deferred here. Beyond the volume, the relocation of the black-hole puzzle invites a precise question: whether the joint structure of the macrostate and unitary projections of an evaporating black hole can be characterized explicitly enough to state the consistency condition the puzzle demands — a question for the bridge between the present framework and the physics of gravitation, and one we do not pretend the present article is equipped to answer.

Acknowledgments

The author thanks colleagues of the Information Physics Institute and the Institute of Integrative and Interdisciplinary Research for sustained discussion of the foundations of the series.

References

[1] J. D. Bekenstein. Universal upper bound on the entropy-to-energy ratio for bounded systems. *Physical Review D*, 23(2):287–298, 1981.

[2] C. H. Bennett. The thermodynamics of computation—a review. *International Journal of Theoretical Physics*, 21(12):905–940, 1982.

[3] A. Bérut, A. Arakelyan, A. Petrosyan, S. Ciliberto, R. Dillenschneider, and E. Lutz. Experimental verification of Landauer's principle linking information and thermodynamics. *Nature*, 483(7388):187–189, 2012.

[4] S. W. Hawking. Particle creation by black holes. *Communications in Mathematical Physics*, 43(3):199–220, 1975.

[5] B. Kriger. The information substrate: differentiation as the ground of information. The Information Substrate Theory (IST), Volume VII, Article 1. Zenodo, 2026.

[6] B. Kriger. The projection theorem: how known information measures descend from the substrate. The Information Substrate Theory (IST), Volume VII, Article 3. Zenodo, 2026.

[7] B. Kriger. *Structural Persistence and Coherence Analysis: A Topological–Dissipative Method Across Physical Scales*. Monograph series, Volume VI. Zenodo, 2026. doi:10.5281/zenodo.20455316.

[8] R. Landauer. Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5(3):183–191, 1961.

[9] S. Lloyd. Computational capacity of the universe. *Physical Review Letters*, 88(23):237901, 2002.

[10] L. Susskind. The world as a hologram. *Journal of Mathematical Physics*, 36(11):6377–6396, 1995.

[11] M. M. Vopson. The mass-energy-information equivalence principle. *AIP Advances*, 9(9):095206, 2019.

[12] M. M. Vopson. The second law of infodynamics and its implications for the simulated universe hypothesis. *AIP Advances*, 12(7):075310, 2022.

[13] J. A. Wheeler. Information, physics, quantum: the search for links. In W. H. Zurek, editor, *Complexity, Entropy, and the Physics of Information*, pages 3–28. Addison-Wesley, 1990.

A A Decision Procedure for Information Claims about the Universe

The diagnoses of this article can be assembled into a short procedure for classifying any claim of the form "the universe *does something* with information." The procedure is offered not as an algorithm but as a checklist that makes the category of a claim explicit.

(P1) *Identify the quantity.* Does the claim invoke an *amount* (bits, states, entropy), a *process* (operations, transmission), a *conserved content* (preservation across time), or a *constituent* (a fundamental unit or substance)? If none, the claim may be about the substrate sense, and Proposition 3.2 applies: the universe needs it in no form.

(P2) *Find the axis.* For an amount, the axis is channel, description, or macrostate; for a process, the time-projection; for a conserved content, a projection under the time-coordinate; for a constituent, whichever projection supplies the unit (channel for a bit, cost for an erasure energy).

(P3) *Test the subject.* Is the claim predicated of the projection or of the substrate? If of the projection, it is well-formed and may be true physics (storage bounds, computational capacity, unitarity, the Landauer cost). If of the substrate, it commits the relevant locality error (Propositions 4.1, 5.1, 6.1, 7.1) and is ill-typed.

(P4) *Relocate, do not resolve, where a tension appears between projections.* If the claim sets two projections against each other (as the black-hole information puzzle sets the macrostate against the unitary projection), classify it as a joint-structure consistency question and resist the temptation to treat the classification as a solution.

Applied to the pictures of the article, the procedure returns the rows of the table in Section 7.2: storage and the holographic bound pass (P3) as projection statements; processing passes as a time-projection statement and fails only when promoted to the whole; conservation passes as a projection statement and the black-hole puzzle is sent to (P4); and the reifications fail (P3) as projection quantities predicated of the substrate, with the mass–energy–information case marked for the empirical assessment of Article 9.

The Projection Theorem: How Known Information Measures Descend from the Substrate

Differentiatedness as the source of which Shannon, Kolmogorov, and the thermodynamic measures are projections along single axes

Boris Kriger[1,2]

[1] Information Physics Institute
boris.kriger@informationphysicsinstitute.net

[2] Institute of Integrative and Interdisciplinary Research, Toronto
boriskriger@interdisciplinary-institute.org

ORCID: https://orcid.org/0009-0001-0034-2903

The Information Substrate Theory (IST) — Volume VII, Article 3
Part I: The Substrate — load-bearing core of the volume

Abstract

The Information Substrate Theory holds that information is not a substance, a force, or a fundamental constituent of the world, but the name of one ineliminable feature of being: that it is differentiated. The undifferentiated is only nothing, and nothing is impossible; therefore being is necessarily differentiated, and *information* names this differentiatedness rather than adding anything to it. This article supplies the result that earns the word *theory*. We formalize the substrate as a differentiation structure and an *axis* as a reduction that retains only the distinctions registered along one coordinate. We then state and prove a Projection Theorem in the only honest form available: for each axis we exhibit a reduction map and show that the known information measure carrying that name is the composition of a single canonical functional with that reduction, where the functional is fixed — up to the axis's native indeterminacy (units, additive constant) — by characterization theorems already established in the literature. Shannon information is derived as the channel-axis projection, forced by the Shannon–Khinchin uniqueness theorem; Kolmogorov complexity as the description-axis projection, fixed up to a constant by the invariance theorem; Boltzmann–Gibbs entropy as the macrostate-axis projection. Effective complexity, Landauer's bound, and Bateson's "difference" are located by axis but are marked, where appropriate, as not admitting derivation in the same sense. The framework's distinguishing, falsifiable-in-principle content is that it predicts relations *between* projections that each projection-level theory must treat as coincidence: why the expected description length of a source equals its entropy, and why effective complexity is non-monotone in disorder where Kolmogorov complexity is monotone. We make this precise by classifying the maps between axes into three structural types — rescaling, kernel, and refinement — and proving a rigidity theorem for each, so that the proportionality of thermodynamic to channel entropy, the agreement of expected complexity with entropy (exact at the rate level, after Brudno), and the effective-complexity dissociation are shown to be the single forced relation of their respective morphism rather than three coincidences. An explicit epistemic-status ledger separates what is proved, what is derived, and what remains conjectured — the latter now narrowed to the completeness of the axis family and of the morphism typology.

Keywords: information substrate; differentiation; Shannon entropy; Kolmogorov complexity; effective complexity; Boltzmann–Gibbs entropy; Landauer bound; characterization theorems; projection; foundations of physics.

Contents

1 Introduction

A reader surveying the measures that go by the name *information* finds not one quantity but a family whose members disagree about almost everything. Shannon entropy is a property of a probability distribution and is maximized by disorder; Kolmogorov complexity is a property of a single object and is also maximized by disorder, yet the two are computed by entirely different means and one of them is uncomputable. Effective complexity, by contrast, is *low* for both perfectly ordered and perfectly random objects and high only in between. Thermodynamic entropy is dimensional, carrying units of energy over temperature, while the others are dimensionless. Landauer's bound attaches an energetic cost to the erasure of a bit, importing a dimensional constant that none of the purely combinatorial measures possess. Bateson's "difference that makes a difference" is not a number at all but a relation to a receiver.

The standard situation is to treat these as distinct tools for distinct jobs, related only by occasional and slightly surprising theorems — for example, that the expected Kolmogorov complexity of objects drawn from a computable source equals the Shannon entropy of that source up to an additive constant [9, 26]. On the standard view such a result is a fortunate bridge between two separately motivated theories. The present article advances a different reading. It is the claim of the Information Substrate Theory (IST) that these measures are not independent theories that happen to make contact, but *projections of a single structure along different axes*. The structure is differentiatedness as such; the axes are the coordinates — channel, description, macrostate, regularity, thermodynamic cost, receiver — along which differentiatedness can be registered; and each named measure is what one obtains by collapsing the substrate onto one axis and applying the canonical functional native to that axis.

This article is the load-bearing core of Volume VII. The foundational descent — that there is something, that it is necessarily differentiated, that *information* is the name of differentiatedness and not an addition to it — is established in the preceding articles of this volume and in prior work of the series [14–16], and is not re-derived here. We take the substrate as given and ask the question that makes IST a theory rather than a classification: can the known information measures be *derived* from the substrate by reduction along an axis, or can they not? Where derivation succeeds we exhibit it; where it does not, we say so. The deliberately modest form of the central theorem reflects a discipline carried throughout the series: a foundation earns its standing by proving less and establishing it, not by asserting more.

We are careful at the outset to disclaim the manifesto register. IST does not hold that "matter is information," that "information is fundamental," or that "the universe is a computer." On the contrary — as the companion article of this volume argues, the universe *needs* no information in any form, because differentiatedness is not a thing the universe contains but a feature it cannot lack. The classification stance toward predecessors is correspondingly generous: each measure is correct within its axis, and the task is to locate it, never to demote it.

2 Logical Structure of the Argument

The argument proceeds in five dependent steps. The reader may find it useful to hold the following map in view before entering the technical development.

Step 1 (Section 3): The substrate and the notion of an axis.
We fix, within the describable, a formal model of differentiatedness — a *differentiation structure* — and define an *axis* as a reduction that retains only those distinctions registered along one coordinate. The central conceptual move is that the substrate itself carries no axis and is therefore in the domain of no axis-bound measure. This step supplies the objects on which everything later operates.

Step 2 (Section 4): The Projection Theorem.
We state the theorem in the form the foundations permit: for each axis there exist a reduction map and a canonical functional, fixed up to the axis's native indeterminacy by an established characterization theorem, whose composition is the named measure. The theorem is a schema; its substance is the individual derivations and the cross-axis relations they induce.

Step 3 (Sections 5–6): The two core derivations.
We carry out the derivation explicitly for the two measures whose reduction is cleanest. Shannon information is the channel-axis projection, forced by the Shannon–Khinchin uniqueness theorem. Kolmogorov complexity is the description-axis projection, fixed up to an additive constant by the invariance theorem. These are the two derivations the volume promises to deliver.

Step 4 (Section 7): Relations among projections.
Here lies the distinguishing content. Because the projections share a domain, relations among them are not coincidences but consequences. We show that averaging the description-axis projection over a source returns the channel-axis projection (the expected-complexity/entropy identity), and that effective complexity differs from Kolmogorov complexity precisely because it retains only the regularity facet of the same differentiatedness, which accounts for its non-monotonicity. We also locate the macrostate, cost, and receiver axes, deriving Boltzmann–Gibbs entropy and marking honestly what the cost and receiver axes do and do not yield.

Step 5 (Section 8): The structure of induced relations.
Here clause (ii) of the Projection Theorem is converted from a slogan into a proved correspondence. We build the lattice of reductions, classify the maps between axes into three morphism types (rescaling, kernel, refinement), and prove a rigidity theorem for each, showing that the three relations of Step 4 are forced — each into a single functional form — by the morphism connecting the axes. This is the falsifiable-in-principle content, and the new mathematical work the framework itself generates.

Step 6 (Section 10): Epistemic status.
We close the technical development with an explicit ledger separating what is proved

(the imported characterization theorems, the cross-axis identities, and the three rigidity theorems), what is derived (the identification of each measure with an axis projection), and what is conjectured (the completeness of the axis list and of the morphism typology).

The dependencies are linear: Step 2 presupposes Step 1; Steps 3 instantiate Step 2; Step 4 presupposes Step 3; Step 5 presupposes Step 4 and supplies the falsifiable-in-principle content; Step 6 audits the whole. The thesis advanced by the totality is that the plurality of information measures is the plurality of *axes*, not of *things measured*, and that their points of contact are the rigidities of the morphisms between axes.

3 The Substrate and the Notion of an Axis

3.1 Background and related work

The idea that information is registered difference is old and well attested. Bateson [2] defined information as a difference that makes a difference to a receiver; Spencer-Brown [21] built a calculus on the primitive act of drawing a distinction; Shannon [19] made the registered difference quantitative by attaching it to a probability distribution over distinguishable messages. The maximum-entropy program of Jaynes [10, 11] read thermodynamic entropy itself as an inference from incomplete distinction among microstates. What these share is the recognition that a measure of information is always a measure *relative to* some specification of which differences count. IST takes this recognition to its root: prior to any specification of which differences count, there is differentiatedness as such, and each specification is the choice of an axis. The present section formalizes "differentiatedness as such" and "axis" so that the later derivations have determinate objects to operate on.

3.2 The differentiation structure

Within the describable, we model the substrate as a structure whose only primitive is distinction.

Definition 3.1 (Differentiation structure)**.** A *differentiation structure* is a pair $\mathcal{S} = (\Omega, \sim)$ where Ω is a nonempty set of configurations and $\sim$ is the identity relation on Ω, regarded as the *finest* distinction: $x \sim y$ iff $x = y$. We require $|\Omega| \geq 2$. The structure carries no metric, no measure, no ordering, and no coordinate; it records only that distinct configurations are distinct.

The requirement $|\Omega| \geq 2$ is the formal residue of the foundational descent. A structure with a single configuration draws no distinction and is, in the sense established in the prior articles, indistinguishable from no structure at all [14]; differentiatedness begins exactly when there is more than one thing to tell apart. We emphasize that Definition 3.1 is *modal*, not temporal. It does not say that Ω comes to be differentiated; it says that being differentiated is what Ω is. No act, agent, event, or moment appears in the definition, and the later prohibition on temporal predicates for the substrate (Article 4 of this volume) is the prohibition on reading any process into $\sim$.

Definition 3.2 (Axis and reduction)**.** An *axis* α on $\mathcal{S}$ is a surjection $\rho_\alpha\colon \mathcal{S} \to \mathcal{S}_\alpha$ onto an *axis-reduced* structure $\mathcal{S}_\alpha$, together with a class Φ_α of admissible functionals on $\mathcal{S}_\alpha$. The map ρ_α is the *reduction along* α: it retains exactly the distinctions registered by the coordinate α and collapses all others. We call $\mathcal{S}_\alpha$ the *α-projection* of the substrate.

Three axis-reduced structures will concern us, corresponding to three coordinates.

- **Channel reduction.** ρ_{ch} forgets the identity of configurations and retains only how often each is distinguished from the rest under repetition: it equips Ω with a probability vector $p = (p_1, \ldots, p_m)$, $\sum_i p_i = 1$, and discards everything else. $\mathcal{S}_{\mathrm{ch}}$ is the simplex of finite probability vectors.
- **Description reduction.** ρ_{de} retains a single configuration x as a finite object to be reconstructed and discards the ensemble: $\mathcal{S}_{\mathrm{de}}$ is the set of finite binary strings, $x \in \{0,1\}^*$.
- **Macrostate reduction.** ρ_{ma} retains only a coarse-graining: a partition of Ω into macrostates and the count W of microstates compatible with the observed macrostate.

3.3 What it means to "derive a measure from the substrate"

We can now state precisely the burden the title of this article accepts.

Definition 3.3 (Axis derivation)**.** A real- or order-valued information measure μ is *derived from the substrate along axis* α if there is a functional $M_\alpha \in \Phi_\alpha$ such that

$$\mu = M_\alpha \circ \rho_\alpha, \tag{1}$$

and M_α is singled out within Φ_α — up to the native indeterminacy of the axis (a positive multiplicative constant fixing units, or an additive constant) — by a characterization theorem stated in conditions intrinsic to $\mathcal{S}_\alpha$.

The phrase "up to the native indeterminacy of the axis" is essential and is not a loophole. A projection cannot recover what the reduction discarded; in particular it cannot recover a choice of units (channel and macrostate axes) or the choice of universal machine (description axis). What a derivation *can* establish is that, once the axis is fixed, the *functional form* of the measure is not a further free choice but is forced by conditions that make no reference to anything outside $\mathcal{S}_\alpha$. This is the strongest sense of derivation available for a quantity that is, by construction, a projection.

3.4 Counterarguments and Responses

Objection 1 (the substrate is just a probability space in disguise). One might suspect that Definition 3.1 secretly smuggles in a measure and that the "substrate" is a measure space under another name. *Response.* It does not: $\mathcal{S}$ carries no measure. The probability vector appears only *after* the channel reduction ρ_{ch}, and a different reduction (description, macrostate) produces an object with no probabilities in it at all. The substrate is precisely what is common to all three reductions and is therefore poorer than any of them, not richer. If it already were a probability space, the description axis could not be a reduction of it without probabilities, which is false.

Objection 2 (the identity relation is trivial and carries no content). *Response.* The triviality is the point. The substrate is meant to carry the minimum that distinguishes being from nothing, no more. All structure beyond bare distinguishability is introduced by an axis and is thereby disclosed as axis-relative rather than intrinsic. A richer primitive would prejudge which axis is fundamental, which is exactly the prejudgment IST refuses.

Objection 3 (modeling the substrate at all violates the boundary). Since later articles insist the substrate is measured by no axis, is it coherent to give it any formal model? *Response.* The model is of differentiatedness *within the describable*, where it functions as the common domain of the reductions. The boundary prohibition concerns predicates that cross out of the describable (a simulator, a creator, temporal predicates applied to the substrate itself); it does not prohibit specifying the domain on which projections act. The distinction is maintained explicitly in Article 4.

3.5 Section Result and Implications

We have fixed three objects: a differentiation structure $\mathcal{S}$ carrying only bare distinguishability; the notion of an axis as a reduction retaining one coordinate's distinctions; and a precise sense in which a measure is *derived* from the substrate along an axis. The implication carried forward is structural: any quantity that depends on probabilities, on description length, or on macrostate counts is by construction a function of an axis-reduced object, not of the substrate, and the substrate is the unique object on which no such quantity is defined. This sets up the theorem of the next section.

4 The Projection Theorem

4.1 Statement

Theorem 4.1 (Projection Theorem)**.** *Let $\mathcal{S}$ be a differentiation structure and let $\mathcal{A} = \{\mathrm{ch}, \mathrm{de}, \mathrm{ma}, \dots\}$ be a family of axes in the sense of Definition 3.2. For each $\alpha \in \mathcal{A}$ for which a characterization theorem of the form required by Definition 3.3 exists, there is* *a reduction ρ_α and a functional M_α, unique within Φ_α up to the native indeterminacy of α, such that the established information measure bearing the name of axis α equals $M_\alpha \circ \rho_\alpha$. Consequently:*

(i) (Locality of measures.) *Each such measure is a function of $\rho_\alpha(\mathcal{S})$ alone and is undefined on $\mathcal{S}$ itself, which lies in the domain of no M_α.*

(ii) (Induced relations.) *Any relation among the reductions $\{\rho_\alpha\}$ induces a relation among the measures $\{M_\alpha \circ \rho_\alpha\}$; such relations are consequences of the shared domain $\mathcal{S}$, not independent coincidences.*

(iii) (Honest incompleteness.) *For axes admitting no such characterization theorem, no derivation in the sense of Definition 3.3 is claimed; the associated quantity is* located *by its axis but not derived.*

Theorem 4.1 is, as stated, a schema. Its content is exhausted by (a) the individual derivations that instantiate it — carried out for ch, de, and ma in Sections 5–7 — and (b) the induced relations of clause (ii), which are the theory's predictive content. We prove the schema by proving its instances; the schema itself asserts only that the instances have the common form (1) and that their relations descend from a shared domain. We resist the temptation to dress the schema as more than this.

Remark 4.2 (Clause (ii) is made precise and proved in Section 8)**.** As stated, clause (ii) asserts only that cross-axis relations are induced. This is not yet falsifiable, because "induced" is undefined. Section 8 supplies the definition — a relation is induced by an inter-axis morphism together with the measures' characterizations — and proves, for each of the three relations of Section 7, a rigidity theorem that forces its functional form. The reader who wishes to see the predictive content of clause (ii) discharged in full should read Section 8 as the completion of the present statement.

4.2 Why a schema and not a single closed-form law

A reader trained in physics may expect a foundational theorem to take the form of a single equation from which the special cases drop out by substitution. No such equation is available here, and pretending otherwise would be the manifesto error the series forbids. The reason is intrinsic: the axes have *different codomains and different native indeterminacies*. The channel and macrostate axes yield additive, unit-bearing functionals; the description axis yields a functional defined only up to an additive constant and uncomputable in general; the receiver axis yields no number at all. A single closed form would have to suppress these differences, which are not defects of presentation but the very facts that distinguish the axes. The honest unification is therefore structural — one domain, many projections — rather than formulaic.

4.3 Counterarguments and Responses

Objection 1 (a schema is not a theorem). If the substance is in the instances, what does Theorem 4.1 add? *Response.* It adds two non-trivial assertions that the instances alone do not: locality (clause i)—that no known measure is a function of the substrate itself—and the status of cross-axis relations (clause ii)—that they are entailed rather than contingent. Both are falsifiable in principle. Locality fails if some bona fide information measure can be shown to depend on distinctions that survive *every* reduction; the induced-relations claim fails if a cross-axis relation is exhibited that provably cannot be obtained from any relation among the reductions. The schema thus has empirical-in-principle bite even though it is not a closed form.

Objection 2 (the axis family is gerrymandered to fit the measures). *Response.* This is a genuine risk and we do not claim the list $\mathcal{A}$ is complete or canonical; clause (iii) and Section 10 mark this explicitly as conjecture. The defense is that each axis is specified by an *independent* reduction with its own characterization theorem drawn from the prior literature, not reverse-engineered from the measure. The channel axis is fixed by the Shannon–Khinchin conditions, formulated long before IST; the description axis by the invariance theorem; the macrostate axis by the microcanonical counting. The axes were not invented to host the measures.

4.4 Section Result and Implications

The Projection Theorem fixes the form of the entire volume's classificatory and derivational work: to locate a measure is to identify its axis and reduction; to derive it is to exhibit the characterization theorem that forces its functional form on that axis; and to predict is to read a relation among measures off a relation among reductions. The remainder of the article discharges this program for the principal measures and then audits the result.

5 The Channel Axis: Deriving Shannon Information

5.1 The reduction and the functional

The channel reduction ρ_{ch} sends $\mathcal{S}$ to a finite probability vector $p = (p_1, \ldots, p_m)$ on the m distinguished configurations. We seek the functional on $\mathcal{S}_{\text{ch}}$ that measures the amount of differentiation registered per draw. The admissible class Φ_{ch} is fixed by conditions intrinsic to a probability vector, and the functional is then not chosen but forced.

Theorem 5.1 (Shannon–Khinchin uniqueness; 6, 12, 19)**.** *Let* $\{H_m\colon \Delta_{m-1} \to \mathbb{R}_{\geq 0}\}_{m\geq 2}$ *be a family of functionals on finite probability vectors satisfying:*

(K1) (Continuity.) H_m *is continuous in* p*.*

(K2) (Maximality.) *For fixed m, H_m is maximized by the uniform vector $u_m = (1/m, \ldots, 1/m)$.*

(K3) (Expansibility.) *Adjoining an impossible outcome leaves the value unchanged: $H_{m+1}(p_1, \ldots, p_m, 0) = H_m(p_1, \ldots, p_m)$.*

(K4) (Grouping / recursivity.) *For a refinement of one outcome,*

$$H\big(p_1, \ldots, p_{m-1}, q_1, q_2\big) = H\big(p_1, \ldots, p_{m-1}, p_m\big) + p_m\, H\Big(\tfrac{q_1}{p_m}, \tfrac{q_2}{p_m}\Big), \quad q_1 + q_2 = p_m.$$

Then there is a constant $\lambda > 0$ with

$$H_m(p) = -\lambda \sum_{i=1}^{m} p_i \log p_i, \tag{2}$$

the choice of λ (equivalently, the base of the logarithm) being the unit.

A proof is standard and we summarize it in Appendix A. The content for IST is read off directly.

Corollary 5.2 (Shannon as channel-axis projection)**.** *Define $M_{\mathrm{ch}} = H_m$ as in (2). Then Shannon information is the channel-axis projection of the substrate, $\mu_{\mathrm{Shannon}} = M_{\mathrm{ch}} \circ \rho_{\mathrm{ch}}$, and M_{ch} is unique within $\Phi_{\mathrm{ch}} = \{$functionals satisfying (K1)–(K4)$\}$ up to the unit λ.*

Proof. $\rho_{\mathrm{ch}}(\mathcal{S}) = p$ by the definition of the channel reduction. Conditions (K1)–(K4) are conditions on $\mathcal{S}_{\mathrm{ch}}$ alone (they mention only probability vectors and their refinements), so they qualify under Definition 3.3. Theorem 5.1 forces the form (2) up to λ. Hence $M_{\mathrm{ch}} \circ \rho_{\mathrm{ch}}$ is Shannon's H, derived in the sense of Definition 3.3. □

5.2 Interpretation

The derivation says something sharper than "Shannon entropy measures information." It says that *once the substrate is reduced to an ensemble of distinguished outcomes with weights*, the only continuous, additive, refinement- respecting measure of registered differentiation is Shannon's, up to the choice of unit. The much-discussed "observer dependence" of Shannon information is, on this reading, exactly the dependence of the *reduction* ρ_{ch} on a prior decision about which outcomes count as distinct and how often — a decision that selects the channel axis. The functional adds no further observer; the observer is already spent in choosing the axis. This is why the series treats "information" as a spectator-relative word in physics while nonetheless admitting Shannon's H as a perfectly determinate functional: the spectator enters with the axis, not with the formula.

5.3 Counterarguments and Responses

Objection 1 (the grouping axiom is exactly additivity, so the result is circular). *Response.* (K4) is a structural requirement — that the measure of a two-stage distinction be the measure of the first stage plus the conditional measure of the second, weighted by the probability of reaching it — and it is independently motivated by the demand that the registered differentiation of a refinement decompose along the refinement. It is not the conclusion (2) in disguise; many non-Shannon functionals satisfy continuity and maximality and fail only (K4). The Tsallis and Rényi families, for instance, satisfy (K1)–(K3) but replace (K4) with a deformed composition rule, and are thereby excluded. That a different grouping rule yields a different functional is not circularity but the statement that the *additive* channel axis is one axis among possible non-additive ones; IST locates Shannon precisely on the additive one.

Objection 2 (Rényi and Tsallis entropies are also "information," so the derivation is incomplete). *Response.* Agreed, and this is a virtue. They correspond to channel-type axes with a deformed composition law, and IST locates them as such rather than denying them the name. The Projection Theorem's clause (ii) predicts that their relations to Shannon's H (the Rényi family converges to Shannon as the order $\to 1$) reflect a one-parameter family of channel reductions collapsing to the additive one — a relation among reductions, exactly as the schema requires.

5.4 Section Result and Implications

Shannon information is derived as the channel-axis projection of the substrate, its functional form forced by the Shannon–Khinchin conditions up to the unit. The implication is that nothing in Shannon's quantity reaches past the channel reduction to the substrate; it is a measure of an ensemble, and the ensemble is already a projection. We now perform the analogous derivation on a reduction that retains a single object rather than an ensemble.

6 The Description Axis: Deriving Kolmogorov Complexity

6.1 The reduction and the functional

The description reduction ρ_{de} retains a single distinguished configuration as a finite binary string x to be reconstructed, discarding the ensemble. The functional native to this axis is description length on a universal machine.

Definition 6.1 (Kolmogorov complexity)**.** For a universal Turing machine U, the *Kolmogorov complexity* of $x \in \{0,1\}^*$ is

$$K_U(x) = \min\{\, |p| : U(p) = x \,\}, \tag{3}$$

the length of the shortest program that outputs x and halts.

Theorem 6.2 (Invariance; 5, 13, 18, 20)**.** *For any two universal machines U, V there is a constant $c_{U,V}$, independent of x, with $|K_U(x) - K_V(x)| \leq c_{U,V}$ for all x. Hence K is well defined up to an additive constant, written $K(x)$.*

Corollary 6.3 (Kolmogorov as description-axis projection)**.** *Define $M_{\mathrm{de}} = K$. Then Kolmogorov complexity is the description-axis projection of the substrate, $\mu_K = M_{\mathrm{de}} \circ \rho_{\mathrm{de}}$, and M_{de} is fixed up to an additive constant — the native indeterminacy of the description axis — by Theorem 6.2.*

Proof. $\rho_{\mathrm{de}}(\mathcal{S}) = x \in \{0,1\}^*$ by the definition of the description reduction. The minimization (3) refers only to $\mathcal{S}_{\mathrm{de}}$ (strings and a machine), and Theorem 6.2 singles out K up to the additive constant $c_{U,V}$, the choice of U being the native indeterminacy here exactly as the unit λ was for the channel axis. Hence $K = M_{\mathrm{de}} \circ \rho_{\mathrm{de}}$ is derived in the sense of Definition 3.3. □

6.2 Interpretation

The two derivations make the axis structure vivid by contrast. The channel reduction keeps the *ensemble* and forgets which configuration is realized; the description reduction keeps the *realized configuration* and forgets the ensemble. They are reductions of the same substrate along orthogonal coordinates: one toward the statistics of repetition, the other toward the specification of an individual. That both are "information" — and both maximized by disorder — is not a verbal accident but a consequence of their common origin in differentiatedness: a maximally differentiated ensemble has maximal channel entropy, and a maximally differentiated single object has maximal description length, because in each case the reduction retains a maximally distinguished object of its own kind.

6.3 Counterarguments and Responses

Objection 1 (K is uncomputable, so it is not a measure one can apply). *Response.* Uncomputability bears on *evaluation*, not on *derivation*. Definition 3.3 requires that the functional be singled out on its axis by a characterization theorem; the invariance theorem does this regardless of computability. The practical consequence — that one must use a computable compressor as an upper proxy — is exactly what we exploit in the empirical illustration of Section 9, and the gap between proxy and K is acknowledged there rather than hidden.

Objection 2 (the additive constant makes K ill-defined for short strings). *Response.* True, and it is the honest analogue of the unit freedom in Shannon's H. Neither axis can recover what its reduction discarded: the channel axis cannot fix a unit, the description axis cannot fix a machine. Derivation establishes the functional form modulo this irreducible freedom, and asymptotic statements (Section 7) are precisely those that survive it.

6.4 Section Result and Implications

Kolmogorov complexity is derived as the description-axis projection of the substrate, fixed up to an additive constant by the invariance theorem. With two orthogonal projections in hand — channel and description — the Projection Theorem's clause (ii) becomes operative: relations between these projections should now be readable as consequences of their shared domain. The next section shows that they are.

7 Relations Among Projections

This section carries the distinguishing content of the theory. A projection-level theory — one that takes a single measure as primitive — can record that two measures happen to be related but cannot explain why; the substrate, being the common domain, can. We treat three relations: the agreement of expected description length with entropy, the dissociation of effective complexity from Kolmogorov complexity, and the placement of thermodynamic entropy. We then locate the cost and receiver axes and mark honestly what they yield.

7.1 Averaging the description axis returns the channel axis

Proposition 7.1 (Expected complexity equals entropy; 9, 26)**.** *Let p be a computable probability mass function on $\{0,1\}^*$ with finite Shannon entropy $H(p)$. Then*

$$0 \;\leq\; \sum_x p(x)\,K(x) \;-\; H(p) \;\leq\; K(p) + O(1), \tag{4}$$

where $K(p)$ is the complexity of the program computing p. In particular, for a fixed source the expected description-axis projection equals the channel-axis projection up to a constant independent of the data.

The proof is standard (lower bound by the Kraft inequality and the no-hypercompression property of the universal distribution; upper bound by the Shannon–Fano code for p, which has expected length $H(p) + O(1)$ and is itself a program) and is recalled in Appendix B. The point for IST is the interpretation.

IST reading. On the standard view, (4) is a bridge theorem relating two separately motivated quantities, and its existence is mildly surprising: why should the average shortest program length for objects from a source equal that source's channel entropy? On IST it is not surprising and not a bridge between separate theories. The two sides are the same projection *evaluated in two orders.* The left side projects to the description axis first (compute $K(x)$ for each x) and averages over the source second; the right side projects to the channel axis (form p) and measures. Equation (4) states that these commute up to the constant cost of specifying the source — which is exactly what one expects when both axes are coordinates on one differentiated object and "average over the source" is an operation internal to the shared domain. The constant $K(p) + O(1)$ is the price of naming the axis along which the averaging is done; it is the same kind of axis-indeterminacy term that appeared as a unit and as a machine constant above. The relation is thus *induced*, in the precise sense of Theorem 4.1(ii).

7.2 Why effective complexity is not Kolmogorov complexity

Gell-Mann and Lloyd [7, 8] introduced effective complexity to capture the length of a concise description of an object's *regularities*, setting aside its incompressible, random part. Formally one separates a description of x into a part specifying an ensemble E of which x is a typical member and a part specifying x within E; the effective complexity $K_{\text{eff}}(x)$ is the length of the first part alone (made precise through the algorithmic minimal sufficient statistic). The empirically central fact is a *dissociation*: K_{eff} is low for both highly ordered and highly random objects and is large only for objects with rich intermediate structure, whereas K increases monotonically with disorder and is maximal for the random object.

Proposition 7.2 (Regularity-facet retention)**.** *K_{eff} is the projection along a* regularity axis ρ_{re} *that retains, from the description-axis object x, only the distinctions that recur across the ensemble of which x is typical, discarding the distinctions specific to x (its random part). Under this reduction:* (i) *a perfectly ordered x has a near-empty recurring-distinction set and hence small K_{eff};* (ii) *a perfectly random x is typical of the maximal ensemble, whose specification is short, and hence also has small K_{eff};* (iii) $K = K_{\text{eff}} +$ (*description of x within its ensemble*)*, so K retains the random facet that K_{eff} discards.*

IST reading and status. Clauses (i)–(ii) are the standard qualitative behavior of effective complexity and are established in the cited literature; clause (iii) is the standard two-part-code decomposition. The IST contribution is the *identification* of the first part with a distinct axis: K_{eff} and K are projections of the *same* description-axis object along whether-or-not the distinction recurs. This is why they behave differently — not because they measure different things in the loose sense, but because they retain different *facets* of one differentiatedness: K retains regular and random distinctions alike, K_{eff} retains only the regular. The non-monotonicity of K_{eff} in disorder, mysterious if effective complexity is taken as primitive, is on IST the direct signature of projecting out the random facet, since both extremes of disorder have an impoverished regular facet. We mark the axis identification as *derived*; the claim that the regularity axis is fully captured by the algorithmic minimal sufficient statistic (rather than by some refinement) we mark as *conjectured*, as the precise individuation of the regular facet remains an open technical question in the effective-complexity literature.

7.3 Boltzmann–Gibbs entropy as the macrostate-axis projection

Proposition 7.3 (Thermodynamic entropy as macrostate projection)**.** *Let ρ_{ma} reduce $\mathcal{S}$ to a macrostate with W compatible, equiprobable microstates. Then the channel functional M_{ch} applied to the uniform vector u_W gives $M_{\mathrm{ch}}(u_W) = \lambda \log W$, and with the unit fixed to Boltzmann's constant, $\lambda = k_B$ (natural log), this is the Boltzmann entropy $S = k_B \ln W$. For non-uniform microstate weights p_i the same functional gives the Gibbs entropy $S = -k_B \sum_i p_i \ln p_i$.*

Proof. Immediate from (2) evaluated on u_W and on a general p, identifying the unit λ with k_B. The maximum-entropy reading of Jaynes [10, 11] is the statement that, absent further constraints, the macrostate reduction selects the uniform (or, under constraints, the maximum-entropy) microstate weights. □

IST reading. Thermodynamic entropy is the same functional as Shannon's, projected onto the macrostate axis and equipped with the dimensional unit k_B. The dimension that distinguishes thermodynamic entropy from the dimensionless channel measure is precisely the unit freedom λ of the channel axis, now fixed by an external physical constant. This is consistent with the series' position that *information* is dimensionless and acquires energetic meaning only through an external dimensional bridge: the bridge is the choice $\lambda = k_B$, and it is the constant, not the bits, that carries the physics.

7.4 Locating the cost and receiver axes

The thermodynamic-cost axis (Landauer). Landauer [17] and Bennett [3] attach to the logically irreversible erasure of one bit a minimum dissipated energy $k_B T \ln 2$. In the present terms this is a projection onto a *cost axis* that retains, from a logical operation on distinctions, only its minimum thermodynamic price, and it is constituted by the same external dimensional bridge $k_B T$ that appeared in Proposition 7.3. We *locate* the Landauer bound on the cost axis but do not claim to *derive* a measure in the sense of Definition 3.3, because the bound is a constrained minimum fixed by physics (the second law together with the dimensional bridge), not a functional singled out by conditions intrinsic to a reduced combinatorial structure. Proposals that extend this bridge into a putative mass–energy–information equivalence or a "second law of infodynamics" [23, 24] make further, stronger claims; their assessment against established physics is the dedicated subject of Article 9 of this volume, and we neither endorse nor refute them here. We note only, in the vocabulary of the present article, that such proposals reify a particular cost-axis projection as a property of the substrate, which is exactly the projection-for-substrate error the Projection Theorem is designed to detect.

The receiver axis (Bateson). Bateson [2] characterized information as a difference that makes a difference to a receiver. This is a genuine axis — the reduction retains, from the substrate, only the distinctions that a specified receiver is organized to register — but it yields no canonical numerical functional, because the reduction is parameterized by the receiver's organization and the relevant prior work establishes no characterization theorem fixing a unique measure on it. We therefore *locate* the Bateson conception on the receiver axis and explicitly do not derive a measure; "derivation" is not the right word for an axis whose codomain is not numerical. This honesty is the content of Theorem 4.1(iii).

7.5 Counterarguments and Responses

Objection 1 (Proposition 7.1 is a known theorem; IST adds only words). *Response.* The theorem is indeed known; what is new is the claim that it is *not* a coincidence but is forced by a shared domain, and the resulting prediction (clause ii) that every robust cross-axis relation should be similarly derivable from a relation among reductions. That prediction is contentful: it forbids robust cross-axis relations that cannot be so obtained, and a single counterexample would refute the framework. A theory can earn its keep by explaining why known results are not accidental, provided it also stakes a falsifiable claim, which here it does.

Objection 2 (effective complexity has well-known definitional ambiguities; the regularity-axis story inherits them). *Response.* It does, and we mark exactly that as conjecture rather than concealing it. The qualitative dissociation (low at both extremes) is robust across reasonable formalizations and is what the regularity-axis reading explains; the precise individuation of the regular facet is open, and the framework's claim is conditional on its eventual resolution. Marking this boundary is the discipline the volume requires.

Objection 3 (placing thermodynamic entropy on the same functional as Shannon's erases a real physical distinction). *Response.* It does not erase it; it locates it in the unit. The physical distinction between a dimensionless count and an energetic entropy is the dimensional bridge $\lambda = k_B$, and the framework names this bridge as the carrier of the physics. Far from collapsing thermodynamics into combinatorics, the reading isolates precisely the non-combinatorial ingredient.

7.6 Section Result and Implications

Two cross-axis relations have been exhibited as consequences of the shared domain — the expected-complexity/entropy identity and the effective-complexity dissociation — and three further measures (Boltzmann–Gibbs, Landauer, Bateson) have been placed, with Boltzmann–Gibbs derived and the cost and receiver axes honestly marked as located but not derived. The implication is that the apparent disunity of the information measures is the disunity of *axes*, while their points of contact are the *visibility, from the shared substrate, of relations invisible to each axis taken alone.* This is the sense in which IST is a theory of the measures rather than one more measure among them.

8 The Structure of Induced Relations

Clause (ii) of Theorem 4.1 was stated in Section 7 as a slogan: cross-axis relations are induced by the shared domain rather than independent coincidences. A slogan is not falsifiable until "induced" is given a definition precise enough that the inducedness of a *particular* relation can be proved or refuted. This section supplies the definition and the proofs. It develops the lattice of reductions over the substrate (§8.1), specifies what makes a relation *robust* (§8.2), classifies the morphisms along which a relation between two axes can be induced into three types (§8.3), and then proves a rigidity theorem for each type (§8.4–8.6), establishing that each of the three relations of Section 7 is not merely consistent with inducedness but *forced* by it, and forced into a single functional form. The content of this section is the framework's own. It is not imported from the characterization theorems of Sections 5–6; it is generated by the picture of measures as projections over a common domain, and it is precisely what a theory that takes one measure as primitive cannot state. With the rigidity theorems in hand, clause (ii) becomes a refined and falsifiable proposition (§8.7), and what remains conjectural is isolated as a sharp statement about the completeness of the morphism typology rather than a slogan.

8.1 The lattice of reductions

A reduction is a quotient of the substrate, and quotients of a fixed object are partially ordered by fineness. We make this explicit.

Definition 8.1 (Refinement order)**.** For reductions $\rho_\alpha\colon \mathcal{S} \to \mathcal{S}_\alpha$ and $\rho_\beta\colon \mathcal{S} \to \mathcal{S}_\beta$, say α is *coarser than* β, written $\alpha \preceq \beta$, if ρ_α factors through ρ_β: there is a map $q_{\beta\alpha}\colon \mathcal{S}_\beta \to \mathcal{S}_\alpha$ with $\rho_\alpha = q_{\beta\alpha} \circ \rho_\beta$. The relation $\preceq$ is a partial order on reductions (modulo isomorphism of the reduced structures); the finest element is the identity reduction of $\mathcal{S}$ and the coarsest is the terminal reduction to a point.

Two reductions need not be comparable: the channel and description axes are not, since neither factors through the other. What always exists is their *joint reduction*.

Definition 8.2 (Joint reduction and joint structure)**.** The *joint reduction* of α and β is $\rho_{\alpha\beta} = (\rho_\alpha, \rho_\beta)\colon \mathcal{S} \to \mathcal{S}_\alpha \times \mathcal{S}_\beta$. Its image $\Gamma_{\alpha\beta} := \rho_{\alpha\beta}(\mathcal{S}) \subseteq \mathcal{S}_\alpha \times \mathcal{S}_\beta$ is the *joint structure*: the set of pairs of reduced values that some single substrate configuration actually realizes.

The joint structure is the exact record of how the two axes are correlated by their common origin. A point $(a, b) \in \mathcal{S}_\alpha \times \mathcal{S}_\beta$ lies in $\Gamma_{\alpha\beta}$ iff there is a substrate configuration projecting to a along α and to b along β; pairs outside $\Gamma_{\alpha\beta}$ are jointly impossible. This furnishes the bare notion of a cross-axis relation.

Definition 8.3 (Cross-axis relation; induced at the structural level)**.** A *cross-axis relation* between α and β is a subset $R \subseteq \mathcal{S}_\alpha \times \mathcal{S}_\beta$ with $\Gamma_{\alpha\beta} \subseteq R$. The *tightest* cross-axis relation is $\Gamma_{\alpha\beta}$ itself. A relation among the measures, i.e. a constraint $F\big(M_\alpha(a), M_\beta(b)\big) = 0$

on $(a, b) \in \mathcal{S}_\alpha \times \mathcal{S}_\beta$, is *structurally induced* if it holds for every $(a, b) \in \Gamma_{\alpha\beta}$; that is, if it is a property of the joint structure rather than an extra postulate.

Structural inducedness is necessary but, by itself, weak: a relation might hold on $\Gamma_{\alpha\beta}$ by accident of a particular unit or a particular finite substrate, without expressing anything about the axes. Two further ingredients remove the accidents and give the strong notion we need: a robustness requirement (§8.2) and the identification of the structural map that carries one axis to the other (§8.3).

8.2 Robust relations

Each axis carries a native indeterminacy — the unit λ for the channel and macrostate axes, the choice of universal machine for the description axis — and a native notion of *independent composition*, the operation $\otimes$ that combines two independent substrates ($p \otimes q$ for product sources, concatenation for strings, $W_1 W_2$ for microstate counts). A relation that survives only for one choice of unit, or only for a single finite instance, tells us about that choice or that instance, not about the axes. We therefore restrict attention to relations that survive both.

Definition 8.4 (Robust cross-axis relation)**.** A cross-axis relation $F = 0$ is *robust* if it is

(R1) *unit-invariant:* it is preserved under each axis's native rescaling (it may be stated only up to the additive or multiplicative constant the axis cannot fix); and

(R2) *extensive:* it is preserved under independent composition $\otimes$ and under the per-unit normalization $\frac{1}{n}(\cdot)$ in the limit $n \to \infty$, so that it is not an artifact of a single finite substrate.

Robustness is the discipline that makes clause (ii) contentful. It excludes relations true only at small size or for a privileged unit, and retains exactly those that express the standing geometry of the axes over the substrate. All three relations of Section 7 are robust in this sense, and the rigidity theorems below show that robustness alone, together with the structural map between the axes, forces each of them.

8.3 Inter-axis morphisms: three types

The joint structure $\Gamma_{\alpha\beta}$ records that the axes are correlated; the *morphism* records *how*. There are exactly three structural ways one reduced object can stand to another over the same substrate, and they exhaust the relations of Section 7.

Definition 8.5 (Inter-axis morphisms)**.** Let α, β be axes over $\mathcal{S}$.

(i) A *rescaling morphism* Resc is an isomorphism $\iota\colon \mathcal{S}_\alpha \xrightarrow{\sim} \mathcal{S}_\beta$ of reduced structures intertwining the composition laws, $\iota(a \otimes_\alpha a') = \iota(a) \otimes_\beta \iota(a')$, compatible with the reductions on the aligned locus where both are functions of one scalar invariant.

(ii) A *kernel morphism* Ker is a Markov kernel $\kappa\colon \mathcal{S}_\alpha \to \mathcal{S}_\beta$ with $\rho_\beta = \kappa \circ \rho_\alpha$ in law: the β-reduction is obtained from the α-reduction by a stochastic step (sampling).

(iii) A *refinement morphism* Ref is a quotient $q\colon \mathcal{S}_\beta \to \mathcal{S}_\alpha$ with $\rho_\alpha = q \circ \rho_\beta$: α is a coarsening of β (Definition 8.1), forgetting a facet of the finer object.

A measure relation is *induced by* a morphism $m \in \{\mathsf{Resc}, \mathsf{Ker}, \mathsf{Ref}\}$ if it is entailed by the compatibility equation of m together with the characterization theorems of M_α, M_β, with no further postulate.

The three morphisms are structurally distinct: an isomorphism preserves information, a kernel adds randomness, a quotient discards it. We now show that each forces a rigidly constrained form on the robust relation it induces, and that the three relations of Section 7 are the three induced forms.

8.4 Rescaling rigidity (Theorem R1)

We begin with a classical lemma, the engine of the rescaling case.

Lemma 8.6 (Logarithmic rigidity of additive monoid maps)**.** *Let $a\colon \mathbb{N}_{\geq 1} \to \mathbb{R}$ satisfy $a(mn) = a(m) + a(n)$ for all m, n and be nondecreasing. Then $a(n) = c \log n$ for some constant $c \geq 0$.*

Proof. This is the multiplicative Cauchy equation under monotonicity; the argument is the one used in Appendix A to fix $A(m) = \lambda \log m$, applied to a. Additivity gives $a(n^k) = k\, a(n)$, hence $a(n) = \log n \cdot \big(a(n)/\log n\big)$, and monotonicity forces the ratio $a(n)/\log n$ to a single constant $c \geq 0$ by the standard squeeze between $a(2^{\lfloor k \log_2 n \rfloor})$ and $a(2^{\lceil k \log_2 n \rceil})$ as $k \to \infty$ [1]. □

Theorem 8.7 (Rescaling rigidity)**.** *Let α, β be additive-monoidal axes whose reductions agree on the cardinality invariant: on the family U_n of n-fold uniform substrates, with $U_m \otimes U_n = U_{mn}$, both M_α and M_β are nondecreasing and additive. Then on the substructure generated by $\{U_n\}$ under $\otimes$,*

$$M_\beta = \frac{\lambda_\beta}{\lambda_\alpha}\, M_\alpha, \tag{5}$$

and this is the only continuous robust relation between the two measures there. The slope $\lambda_\beta/\lambda_\alpha$ is the ratio of the axes' units and is fixed entirely by the reductions; the relation is induced by the rescaling morphism.

Proof. By additivity and monotonicity on $\{U_n\}$, Lemma 8.6 gives $M_\alpha(U_n) = \lambda_\alpha \log n$ and $M_\beta(U_n) = \lambda_\beta \log n$ for constants $\lambda_\alpha, \lambda_\beta \geq 0$. Eliminating $\log n$ yields $M_\beta(U_n) = (\lambda_\beta/\lambda_\alpha) M_\alpha(U_n)$ for every n. The values $\{(\lambda_\alpha \log n, \lambda_\beta \log n)\}_{n\geq 1}$ are dense in the ray they span; any continuous relation $F(M_\alpha, M_\beta) = 0$ vanishing on a dense subset of the ray vanishes on its closure, the line (5). Unit-invariance (R1) permits only the freedom in $\lambda_\alpha, \lambda_\beta$ already present in the slope; extensivity (R2) is automatic because both sides are additive under $\otimes$. Hence (5) is the unique robust relation, and its slope is a datum of the reductions. □

Corollary 8.8 (Shannon–Boltzmann is the forced rescaling relation)**.** *With $M_{\mathrm{ch}} = H$ in nats ($\lambda_{\mathrm{ch}} = 1$) and $M_{\mathrm{ma}} = S$ with $\lambda_{\mathrm{ma}} = k_B$ (Proposition 7.3), Theorem 8.7 gives $S = k_B H$, i.e. on the microcanonical locus $S = k_B \ln W$. The celebrated proportionality of thermodynamic entropy to an information measure is therefore not a discovery about two independent quantities but the unique robust relation forced by the rescaling morphism between the macrostate and channel axes, with slope equal to the dimensional bridge k_B.*

Falsifiable corollary. If two measures were each additive on the shared composition and aligned on cardinality, yet displayed a *robust nonlinear* relation, Theorem 8.7 would be contradicted: one of them would fail additivity, hence fail to be an additive-monoidal projection of the shared substrate. A robust nonlinear H–S relation would thus refute the placement of thermodynamic entropy on a rescaling morphism of the channel axis. None is known; the framework stakes that none exists.

8.5 Averaging rigidity (Theorem R2)

The channel and description axes are not isomorphic, so no rescaling morphism connects them. They are connected by a kernel: the description object is a *sample* from the channel object. Throughout this subsection K denotes the prefix (self-delimiting) Kolmogorov complexity, for which the Kraft inequality $\sum_x 2^{-K(x)} \leq 1$ holds; this is the standard setting of the expected-complexity results [18].

Lemma 8.9 (Exact lower bound)**.** *For any computable source p, $\sum_x p(x)K(x) \geq H(p)$.*

Proof. By Kraft, the lengths $\{K(x)\}$ are the codeword lengths of a prefix code (after at most a normalization absorbed into the constant). Shannon's noiseless source coding theorem states that the expected length of any prefix code is at least the entropy: $\sum_x p(x)K(x) \geq -\sum_x p(x) \log p(x) = H(p)$. □

Theorem 8.10 (Averaging rigidity)**.** *Let the kernel morphism be the sampling kernel κ_p ("draw $x \sim p$"), under which the description reduction factors through the channel reduction. Then:*

(a) (Finite form, induced and bounded.) *For every computable p,*

$$0 \;\leq\; \mathbb{E}_{x\sim p}[K(x)] - H(p) \;\leq\; K(p) + O(1), \tag{6}$$

where the lower bound is Lemma 8.9 and the upper bound uses the Shannon–Fano code for p, itself a program of length $K(p) + O(1)$. The induced relation is $\mathbb{E}_{x\sim p}[K(x)] \doteq H(p)$, and its slack is bounded by the complexity of specifying the kernel — the cost of naming the source.

(b) (Tightness.) *Both bounds are achieved up to constants: the lower bound for sources of low program-complexity, the upper bound for sources whose description complexity $K(p)$ is large, so no smaller robust gap than the naming cost is possible in general.*

(c) (Rate form, exact.) *For a stationary ergodic source with entropy rate h,*

$$\lim_{n\to\infty} \frac{1}{n} K(x_{1:n}) = h \quad \textit{almost surely and in mean} \tag{7}$$

[4]. Under the per-unit normalization required by robustness (R2), the slack of (6) is divided by n and vanishes; the induced relation between the normalized description-axis and channel-axis projections is exact, with h the unique constant satisfying (7) and equal to the Shannon entropy rate.

Proof. (a) Lower bound: Lemma 8.9. Upper bound: the Shannon–Fano–Elias code built from p assigns x a description of length $\lceil -\log p(x)\rceil + 1$; prefixed by a program for p of length $K(p)+O(1)$ this is a self-delimiting description of x, so $K(x) \le K(p) - \log p(x) + O(1)$, and taking $\mathbb{E}_p$ gives $\mathbb{E}_p[K] \le H(p) + K(p) + O(1)$. (b) The lower extreme is attained when $K(p) = O(1)$ (a simply describable source), reducing (6) to $\mathbb{E}_p[K] = H(p) + O(1)$; the upper extreme follows by choosing a family of sources with $K(p) \to \infty$ and bounded entropy, e.g. point masses on incompressible strings, for which $\mathbb{E}_p[K] = K(x^*) + O(1)$ with x^* of high complexity while $H(p) = 0$. (c) Equation (7) is Brudno's theorem; the uniqueness of h is immediate, and its identification with the Shannon entropy rate is the Shannon–McMillan–Breiman theorem. Normalizing (6) by n and using $K(p_n) = o(n)$ for the n-fold product source sends the slack to zero. □

Interpretation. The expected-complexity identity (Proposition 7.1) is exhibited as the induced relation of the sampling kernel. Its mild air of coincidence on the standard view — why should averaging shortest programs return channel entropy? — dissolves: averaging over the source is the kernel morphism, and the relation is what that morphism forces. The much-discussed $O(1)$ slack is not looseness but the irreducible cost of naming the connecting reduction, the exact counterpart of the unit freedom of R1; and under the normalization that robustness demands, the slack vanishes and the induced relation is *exact.*

Falsifiable corollary. The framework forbids a robust source-averaged relation between K and H with normalized gap bounded away from zero, or with finite gap provably exceeding the naming cost $K(p) + O(1)$. Either would refute the identification of the channel–description connection with the sampling kernel.

8.6 Refinement rigidity (Theorem R3)

Effective complexity sits on neither a rescaling nor a kernel morphism from the description axis but on a refinement morphism: the regularity reduction is a *coarsening* of the description object, retaining the recurring facet and discarding the random one. The relevant decomposition is the algorithmic minimal sufficient statistic [22]: a description of x splits into a model part specifying an ensemble E^* of which x is typical, of length $K_{\text{eff}}(x)$, and a data part locating x within E^*.

Theorem 8.11 (Refinement rigidity)**.** *Let q be the coarsening that sends the description object x to its minimal sufficient statistic E^*, so that $\rho_{\text{re}} = q \circ \rho_{\text{de}}$. Then:*

(a) (Induced additive decomposition.) $K(x) = K_{\text{eff}}(x) + R(x) + O(\log|x|)$*, where $R(x) \geq 0$ is the within-ensemble (random-facet) term; this decomposition is the relation induced by the refinement morphism q.*

(b) (No rescaling relation: negative rigidity.) *There is no constant c and no additive constant with $K(x) = c\,K_{\text{eff}}(x) + O(1)$ for all x. Indeed for incompressible strings of length n, $K(x) = n + O(\log n)$ while $K_{\text{eff}}(x) = O(\log n)$, so $K(x)/K_{\text{eff}}(x) \to \infty$.*

(c) (Forced non-monotonicity.) $K_{\text{eff}} = K - R$ *is small at both extremes of disorder — for ordered x because K itself is small, for random x because $R \approx K$ — and large only for intermediate structure. The non-monotonicity of effective complexity in disorder is thus a forced consequence of projecting out the random facet.*

Proof. (a) is the two-part-code form of the minimal sufficient statistic [22], with $R(x) = K(x \mid E^*) + O(\log|x|)$ the log-cardinality of the optimal ensemble. (b) By a counting argument there are at most $2^n - 1$ programs shorter than n bits, so most strings of length n have $K(x) \geq n$; such a string is typical only of the maximal ensemble $\{0,1\}^n$, whose specification has length $K_{\text{eff}}(x) = K(\text{“all } n\text{-bit strings”}) = O(\log n)$. Hence the ratio diverges and no proportionality with bounded error can hold. (c) For ordered x, $K(x) = O(\log n)$ bounds $K_{\text{eff}} \leq K$; for incompressible x, (b) gives $K_{\text{eff}} = O(\log n)$ while $R \approx n$; for x with genuine intermediate regularities the model part is substantial. The intermediate maximum follows. □

Interpretation and the cross-type prediction. R3 does more than reproduce the known dissociation of effective complexity. It explains the dissociation as the signature of a *morphism type*: because K_{eff} and K are related by a refinement (a coarsening that discards the random facet), they cannot be related by a rescaling, and a proportional law between them is therefore not merely absent but *forbidden.* The framework thus predicts the *form* of a cross-axis relation from the type of morphism connecting the axes: rescaling forces proportionality (R1), kernel forces averaging-equality up to naming cost (R2), refinement forces an additive decomposition that excludes proportionality (R3).

Falsifiable corollary. A robust proportional law $K \propto K_{\text{eff}}$ would refute the refinement placement of effective complexity; conversely, a robust *additive decomposition* of the

Shannon–Boltzmann pair (rather than the proportionality of R1) would indicate they are not related by a rescaling after all. The morphism type and the relation form are tied, and each can falsify the other.

8.7 The refined clause (ii) and what remains conjectural

The three rigidity theorems convert clause (ii) of Theorem 4.1 from a slogan into a proved correspondence over the relations actually at issue.

Theorem 8.12 (Refined inducedness)**.** *Each of the three robust cross-axis relations established in Section 7 is induced by exactly one inter-axis morphism type and is rigidly constrained to the corresponding functional form:*

Axis pair	***Morphism***	***Forced form***
channel – macrostate	Resc *(R1)*	*proportional,* $S = k_B H$
channel – description	Ker *(R2)*	*averaging-equal, exact per unit*
description – regularity	Ref *(R3)*	*additive split, no proportion*

Each entry is a theorem (Corollary 8.8, Theorem 8.10, Theorem 8.11); none is an additional postulate beyond the morphism and the measures' characterizations.

What remains open is no longer whether the relations are induced — they are proved to be — but whether the morphism typology is complete.

> **Morphism-Completeness Conjecture.** Every robust cross-axis relation between projection measures of the substrate is induced by a composite of the three morphism types Resc, Ker, Ref.

This is now a precise and falsifiable statement, in sharp contrast to the slogan it replaces. It is refuted by exhibiting a single robust cross-axis relation that provably cannot be obtained as such a composite; it is supported by every case so far examined, in which the relation is induced by one of the three types alone. We mark the conjecture as conjecture in the ledger of Section 10, but the content it now carries is the content clause (ii) was always meant to have.

8.8 Counterarguments and Responses

Objection 1 (the rigidity theorems are repackaged classical results). Lemma 8.6 is Cauchy's, Theorem 8.10(c) is Brudno's, the decomposition in R3 is the minimal sufficient statistic. *Response.* The imported items are the *inputs*; the contribution is the classification of inter-axis morphisms and the demonstration that each morphism type *forces* a single relation form, so that the three famous relations are not three coincidences but three instances of one structural law. No classical source states that proportionality, averaging-equality, and additive-decomposition are the rigidities of rescaling, kernel, and refinement morphisms over a common domain; that statement, and its falsifiable completeness conjecture, are the framework's. The novelty is not a new theorem of information theory but a theorem *about* information measures, available only once they are seen as projections.

Objection 2 (the aligned-locus restriction in R1 is a cheat). R1 proves proportionality only on the substructure generated by the uniform family. *Response.* The restriction is principled, not evasive. Off the aligned locus the two measures depend on data their reductions do not share (the full probability vector versus the macrostate count), and there is then no robust relation to force — which is correct, since thermodynamic entropy and channel entropy do diverge off the microcanonical locus. R1 forces the relation exactly where a robust relation exists and is silent exactly where none does; the locus marks the boundary of the induced relation rather than a gap in the proof.

Objection 3 (three types look as gerrymandered as the axis list did). *Response.* The three types are not chosen to fit the relations; they are the three structural possibilities for a map between quotients of a common object — isomorphism, stochastic extension, deterministic coarsening — corresponding to preserving, adding, and removing distinction. Their exhaustiveness as *primitive* morphisms is a clean structural claim; what is genuinely open, and flagged as such, is only whether every robust relation reduces to a composite of them, which is the Morphism-Completeness Conjecture.

8.9 Section Result and Implications

Clause (ii) of the Projection Theorem is now a result, not a slogan. We defined the lattice of reductions and the joint structure, isolated the robust relations as the ones that express the standing geometry of the axes, classified the inter-axis morphisms into rescaling, kernel, and refinement, and proved a rigidity theorem for each: proportionality with unit-ratio slope (R1, forcing $S = k_B H$), averaging equality exact under normalization (R2, forcing the Zvonkin–Levin/Brudno identity), and additive decomposition without proportionality (R3, forcing the effective-complexity dissociation). The three relations of Section 7 are thereby shown to be the three induced forms, each rigidly determined by its morphism type. The remaining ambition — that these three types exhaust the ways a robust relation can arise — is stated as the precise, falsifiable Morphism-Completeness Conjecture. This is the sense, promised but not delivered in Section 7, in which the substrate *sees* the relations among projections that each projection conceals: it sees them as the rigidities of the morphisms between reductions of one differentiated whole.

9 Empirical Illustration: The Operational Face of the Identity $E[K] \simeq H$

Proposition 7.1 is exact but invokes the uncomputable K. Its operational shadow is testable with a real compressor used as an upper proxy for K, by the source-coding theorem: for a stationary source a good lossless compressor approaches the entropy rate, so the average compressed length per symbol of sequences drawn from a Bernoulli(θ) source should track the binary entropy $H(\theta) = -\theta \log_2 \theta - (1-\theta) \log_2 (1-\theta)$, descending toward it as the sequence length grows. This is the per-symbol, rate-level face of the averaging rigidity of Theorem 8.10: under normalization the slack of (6) vanishes (Brudno's theorem, Eq. (7)), so the description-axis proxy and the channel-axis value should coincide in the limit.

Figure 1 shows this for a standard general-purpose compressor (bz2) applied to bit-packed i.i.d. Bernoulli sequences. Panel (a) plots, across the source parameter θ, the entropy curve (the channel-axis projection) against the average compressed bits per symbol at $n = 10^6$ (the description-axis proxy); the proxy tracks the curve across the full range. Panel (b) fixes $\theta = 0.20$ ($H = 0.7219$ bits) and shows the proxy decreasing from 1.376 bits per symbol at $n = 10^3$ toward 0.828 bits per symbol at $n = 10^6$ as the sequence lengthens. The residual offset of order 0.1 bit is expected and is reported, not hidden: a general-purpose compressor is an *upper* bound on K, carries block and dictionary overhead, and is not an ideal entropy coder; the underlying relation (4) holds only up to the additive $K(p) + O(1)$, which the proxy realizes as a positive gap. The point of the figure is the direction and robustness of the convergence, which is the operational signature that averaging the description-axis projection over the source returns the channel-axis projection. The generating code is given in Appendix C.

10 Epistemic Status: Proved, Derived, Conjectured

In the manner of the series, we close the technical development with an explicit ledger. Three statuses are distinguished. *Proved* marks results established by theorems imported from the cited literature, whose truth does not depend on IST. *Derived* marks the IST identifications — that a given measure is the projection of the substrate along a named axis — which are sound given the imported theorems and the definitions of Section 3. *Conjectured* marks claims the framework advances but has not established.

Claim	**Status**	**Rests on**
Shannon–Khinchin uniqueness of H	Proved	Faddeev [6], Khinchin [12]
Invariance theorem for K	Proved	Kolmogorov [13], Li and Vitányi [18]
$\sum_x p(x)K(x) = H(p) + O(1)$	Proved	Grünwald and Vitányi [9], Zvonkin and Levin [26]

continued on next page

Epistemic-status ledger (continued)

Claim	Status	Rests on
Effective-complexity dissociation (low at both extremes)	Proved	Gell-Mann and Lloyd [7, 8]
Per-symbol K → entropy rate (ergodic sources)	Proved	Brudno [4]
Algorithmic minimal sufficient statistic decomposition	Proved	Vereshchagin and Vitányi [22]
Shannon = channel-axis projection	Derived	Cor. 5.2
Kolmogorov = description-axis projection	Derived	Cor. 6.3
Boltzmann–Gibbs = macrostate-axis projection	Derived	Prop. 7.3
K_{eff} = regularity-axis projection of x	Derived	Prop. 7.2
Rescaling rigidity: $\mathsf{Resc} \Rightarrow S = k_B H$	Proved (IST framing)	Thm. 8.7, Cor. 8.8
Averaging rigidity: $\mathsf{Ker} \Rightarrow \mathbb{E}[K] \doteq H$, exact per unit	Proved (IST framing)	Thm. 8.10
Refinement rigidity: Ref ⇒ split, no proportion	Proved (IST framing)	Thm. 8.11
Each §7 relation induced by one morphism type	Derived	Thm. 8.12
Landauer bound located on a cost axis	Derived (located)	§7.4
Bateson "difference" on a receiver axis	Located only	§7.4
Regular facet = algorithmic minimal sufficient statistic exactly	Conjectured	open in cited literature
The axis family $\mathcal{A}$ is complete / canonical	Conjectured	§4.3
Morphism typology is complete (every robust relation a composite)	Conjectured (falsifiable)	§8.7

The ledger is the article's guard against the manifesto register. The load-bearing derivations — Shannon and Kolmogorov as projections — are sound, and the inducedness of each cross-axis relation of Section 7 is now *proved* by the rigidity theorems of Section 8 rather than asserted. What remains conjectural has correspondingly narrowed: not whether the relations are induced, but whether the axis family and the morphism typology are complete. Both residual conjectures are precise and falsifiable, and neither is asserted as established.

11 Discussion

The result of this article is narrow by design and consequential in proportion. It does not announce a new ontology, place information among the fundamental constituents, or rank the substrate with the programs of Wheeler [25] and his successors. It makes the smaller claim that the recognized measures of information are projections of one structure along different axes, and it backs the claim by deriving the two principal mea-

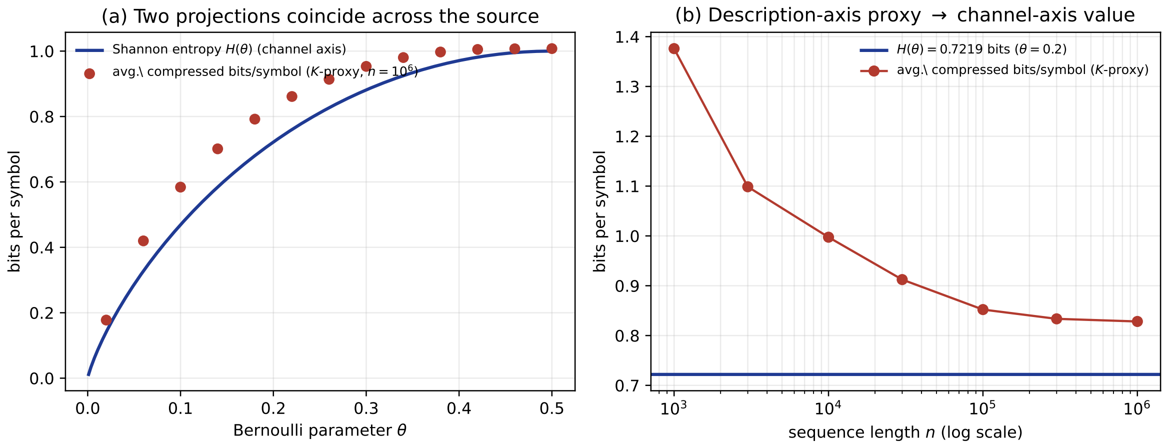

Figure 1: The operational face of Proposition 7.1. **(a)** The Shannon entropy $H(\theta)$ (channel-axis projection, solid) and the average compressed bits per symbol (a computable upper proxy for the description-axis projection K, points) coincide across the source parameter θ at $n = 10^6$. **(b)** At fixed $\theta = 0.20$ the proxy descends toward $H(\theta) = 0.7219$ bits as the sequence length n increases. The residual gap is the compressor's overhead and the additive constant of (4), and is reported rather than suppressed. Values are illustrative of the trend, not a precise estimate of K, which is uncomputable.

sures from explicit reductions using characterization theorems that long predate the framework. The payoff is the predictive content of the Projection Theorem's second clause, which Section 8 discharges in full: relations among the measures that each measure, taken as primitive, can only record as brute fact become, on the substrate reading, the rigidities of the morphisms between reductions of one differentiated whole. We did not merely indicate that the three relations are induced; we proved it, classifying the inter-axis morphisms into rescaling, kernel, and refinement, and showing that each type forces a single functional form — proportionality $S = k_B H$, averaging-equality exact under normalization, and additive decomposition without proportionality, respectively. This also answers the natural worry that IST merely repackages imported uniqueness theorems. The characterization theorems of Khinchin, Kolmogorov, and Brudno are inputs; the theorem that proportionality, averaging-equality, and additive-decomposition are the rigidities of three structural morphism types over a common domain — and the falsifiable conjecture that these three exhaust the robust relations — are statements about information measures that only the substrate picture makes available, and that no projection-level theory can formulate.

The reading also clarifies a tension that has accompanied the series. The prior volumes argued that *information* is a spectator-relative, dimensionless, and multiply defined word that should not be admitted as a fundamental physical concept, and recommended *constraint* in its place. The present volume is titled for information. There is no reversal. IST does not promote information to a fundamental concept; it identifies information with differentiatedness, which is not a thing the world contains but a feature it cannot lack, and then shows that every *quantitative* notion of information is a projection requiring a prior choice of axis — the very spectator-relativity the earlier volumes flagged, now exhibited as the choice of ρ_α. The dimensionality of thermodynamic entropy is isolated as an external bridge ($\lambda = k_B$), confirming rather than contradicting the claim that the bits themselves are dimensionless. The volume thus continues the discipline of the series: it says less than the manifesto and proves what it says.

11.1 Limitations

Four limitations are material. *First*, the axis family $\mathcal{A}$ is not shown to be complete or canonical; the framework would be strengthened by an independent principle generating the axes rather than a curated list, and weakened by the discovery of a bona fide measure fitting no axis. *Second*, the regularity-axis identification inherits the unresolved definitional questions of effective complexity, and is conditional on their resolution. *Third*, the central theorem is a schema whose content is its instances; it is not a closed-form law, and readers seeking a single generating equation will find the structural unification less satisfying than a formulaic one — though we have argued the structural form is the honest one. *Fourth*, the empirical illustration of Section 9 uses a computable compressor as a proxy for an uncomputable quantity, and exhibits a trend rather than a precise estimate; it corroborates the direction of Proposition 7.1 but cannot, in principle, measure K itself.

11.2 Future Directions

Three directions follow. The first is the search for an axis-generating principle: if the admissible reductions of a differentiation structure can be classified intrinsically (for instance, by the algebraic type of the equivalence the reduction respects), the axis-completeness conjecture in the ledger would become a theorem or be refuted. The second is the analogous question one level up, now that the induced-relations claim has been made precise and proved in the three cases at hand (Section 8): to settle the Morphism-Completeness Conjecture by either deriving the three morphism types — rescaling, kernel, refinement — as the exhaustive primitive maps between quotients of a common object, or exhibiting a robust cross-axis relation induced by none of their composites. This would convert the last remaining slogan of clause (ii) into a theorem. The third is the worked case promised to Article 9: the thermodynamic projection applied to a physical dataset, where the cosmic-web analysis of DESI DR1 furnishes a concrete instance of mistaking a projection (the thermodynamic-cost axis) for the substrate, and where the apparatus of the present article supplies the diagnosis. These directions are pursued in the later articles of the volume; the present article supplies the foundation on which they stand.

Acknowledgments

The author thanks colleagues of the Information Physics Institute and the Institute of Integrative and Interdisciplinary Research for sustained discussion of the foundations of the series.

References

[1] J. Aczél and Z. Daróczy. *On Measures of Information and Their Characterizations.* Academic Press, New York, 1975.

[2] G. Bateson. *Steps to an Ecology of Mind.* Chandler Publishing, San Francisco, 1972.

[3] C. H. Bennett. The thermodynamics of computation—a review. *International Journal of Theoretical Physics*, 21(12):905–940, 1982.

[4] A. A. Brudno. Entropy and the complexity of the trajectories of a dynamical system. *Transactions of the Moscow Mathematical Society*, 44:127–151, 1983.

[5] G. J. Chaitin. On the length of programs for computing finite binary sequences. *Journal of the ACM*, 16(3):407–422, 1969.

[6] D. K. Faddeev. On the concept of entropy of a finite probabilistic scheme. *Uspekhi Matematicheskikh Nauk*, 11(1):227–231, 1956.

[7] M. Gell-Mann and S. Lloyd. Information measures, effective complexity, and total information. *Complexity*, 2(1):44–52, 1996.

[8] M. Gell-Mann and S. Lloyd. Effective complexity. In M. Gell-Mann and C. Tsallis, editors, *Nonextensive Entropy: Interdisciplinary Applications*, pages 387–398. Oxford University Press, 2004.

[9] P. Grünwald and P. Vitányi. Algorithmic information theory. In P. Adriaans and J. van Benthem, editors, *Philosophy of Information*, pages 281–320. North-Holland, Amsterdam, 2008.

[10] E. T. Jaynes. Information theory and statistical mechanics. *Physical Review*, 106(4):620–630, 1957.

[11] E. T. Jaynes. Information theory and statistical mechanics. II. *Physical Review*, 108(2):171–190, 1957.

[12] A. I. Khinchin. *Mathematical Foundations of Information Theory.* Dover Publications, New York, 1957.

[13] A. N. Kolmogorov. Three approaches to the quantitative definition of information. *Problems of Information Transmission*, 1(1):1–7, 1965.

[14] B. Kriger. Differentiation as the ontological condition of actualization. Zenodo, 2026. doi:10.5281/zenodo.18268520.

[15] B. Kriger. Testing nothingness and evaluating the inevitability of life: a deductive chain from being to persistence, intelligence, and the limits of inquiry. Institute of Integrative and Interdisciplinary Research. Zenodo, 2026. doi:10.5281/zenodo.18764754.

[16] B. Kriger. *Structural Persistence and Coherence Analysis: A Topological–Dissipative Method Across Physical Scales.* Monograph series, Volume VI. Zenodo, 2026. doi:10.5281/zenodo.20455316.

[17] R. Landauer. Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5(3):183–191, 1961.

[18] M. Li and P. Vitányi. *An Introduction to Kolmogorov Complexity and Its Appli-*

cations. Springer, 4th edition, 2019.

[19] C. E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423, 1948.

[20] R. J. Solomonoff. A formal theory of inductive inference. Part I. *Information and Control*, 7(1):1–22, 1964.

[21] G. Spencer-Brown. *Laws of Form*. Allen & Unwin, London, 1969.

[22] N. K. Vereshchagin and P. M. B. Vitányi. Kolmogorov's structure functions and model selection. *IEEE Transactions on Information Theory*, 50(12):3265–3290, 2004.

[23] M. M. Vopson. The mass-energy-information equivalence principle. *AIP Advances*, 9(9):095206, 2019.

[24] M. M. Vopson. The second law of infodynamics and its implications for the simulated universe hypothesis. *AIP Advances*, 12(7):075310, 2022.

[25] J. A. Wheeler. Information, physics, quantum: the search for links. In W. H. Zurek, editor, *Complexity, Entropy, and the Physics of Information*, pages 3–28. Addison-Wesley, 1990.

[26] A. K. Zvonkin and L. A. Levin. The complexity of finite objects and the development of the concepts of information and randomness by means of the theory of algorithms. *Russian Mathematical Surveys*, 25(6):83–124, 1970.

A Proof of the Shannon–Khinchin Uniqueness Theorem

We recall the classical argument fixing (2); full treatments are in Khinchin [12] and Aczél and Daróczy [1]. Write $A(m) = H_m(u_m)$ for the value on the uniform vector of size m. Applying the grouping axiom (K4) repeatedly to decompose a uniform distribution on mn outcomes into m groups of n uniform outcomes gives the functional equation

$$A(mn) = A(m) + A(n).$$

By (K2) the value $A(m)$ is non-decreasing in m (a uniform distribution on more outcomes is at least as differentiated), and continuity (K1) extends the Cauchy-type equation to its only non-decreasing solutions, $A(m) = \lambda \log m$ with $\lambda > 0$. For a rational distribution $p_i = k_i/N$ with $\sum_i k_i = N$, group N uniform outcomes into m groups of sizes k_i; (K4) gives

$$A(N) = H_m(p) + \sum_{i=1}^{m} p_i\, A(k_i),$$

whence

$$H_m(p) = \lambda \log N - \lambda \sum_i p_i \log k_i = -\lambda \sum_i p_i \log \frac{k_i}{N} = -\lambda \sum_i p_i \log p_i.$$

Continuity (K1) extends the identity from rational to all p, and expansibility (K3) guarantees consistency under adjoining zero-probability outcomes. This is (2); the constant λ (equivalently the logarithm base) is the unit and is the native indeterminacy of the channel axis. □

B Proof Sketch of Proposition 7.1

Lower bound. The prefix Kolmogorov complexity satisfies the Kraft inequality $\sum_x 2^{-K(x)} \le 1$, so $2^{-K(\cdot)}$ is (up to normalization) a sub-probability and, being lower-semicomputable and dominating every computable measure up to a multiplicative constant, gives $K(x) \le -\log p(x) + K(p) + O(1)$; taking expectations under p and using $\sum_x -p(x)\log p(x) = H(p)$ yields $\sum_x p(x)K(x) \le H(p) + K(p) + O(1)$. *Upper bound for the other side of* (4). The Shannon–Fano–Elias code for p assigns x a codeword of length $\lceil -\log p(x) \rceil + 1$; since this code is itself a program of length $K(p) + O(1)$, no shorter description can have smaller expected length than $H(p)$, giving $\sum_x p(x)K(x) \ge H(p) - O(1)$. Combining, $\left|\sum_x p(x)K(x) - H(p)\right| \le K(p) + O(1)$. Full detail, including the prefix-freeness bookkeeping, is in Li and Vitányi [18], §8, and Grünwald and Vitányi [9]. □

C Code for Figure 1

The following Python script generates Figure 1. It uses a real lossless compressor (bz2) on bit-packed i.i.d. Bernoulli sequences as a computable upper proxy for Kolmogorov complexity, and compares the average compressed length per symbol with the binary Shannon entropy.

```
import bz2
import numpy as np
rng = np.random.default_rng(20260530)

def binary_entropy(theta):
    t = np.clip(theta, 1e-12, 1 - 1e-12)
    return -(t*np.log2(t) + (1-t)*np.log2(1-t))

def packed_bits(bits):
    pad = (-len(bits)) % 8
    if pad:
        bits = np.concatenate([bits, np.zeros(pad, dtype=np.uint8)])
    return np.packbits(bits).tobytes()

def bits_per_symbol(theta, n, reps=1):
    vals = []
    for _ in range(reps):
        bits = (rng.random(n) < theta).astype(np.uint8)
        comp = bz2.compress(packed_bits(bits), 9)
        vals.append(8.0 * len(comp) / n)
    return float(np.mean(vals))

# Panel (a): bits/symbol vs theta at n = 1e6
thetas = np.linspace(0.02, 0.5, 13)
proxy_a = [bits_per_symbol(t, 1_000_000, reps=2) for t in thetas]

# Panel (b): convergence vs n at theta = 0.20
theta_b = 0.20
ns = [1_000, 3_000, 10_000, 30_000, 100_000, 300_000, 1_000_000]
proxy_b = [bits_per_symbol(theta_b, n, reps=4) for n in ns]
H_b = binary_entropy(theta_b)   # = 0.7219 bits
```

The Categorial Boundary: On Predicates That Do Not Cross the Block

The boundary of the describable as the boundary of predicate applicability, and the discipline of an apophasis without negation

Boris Kriger[1,2]

[1] Information Physics Institute
boris.kriger@informationphysicsinstitute.net

[2] Institute of Integrative and Interdisciplinary Research, Toronto
boriskriger@interdisciplinary-institute.org

ORCID: https://orcid.org/0009-0001-0034-2903

The Information Substrate Theory (IST) — Volume VII, Article 4
Part I: The Substrate — foundations of the volume

Abstract

The preceding articles of this volume hold that within the describable the universe is a completed, invariant structure — a block — and that beyond the describable the theory is silent, silent not only of entities but of categories. This article makes that boundary precise and shows what kind of limit it is. It is not a wall in space with a further region behind it, nor a horizon of present ignorance, nor a mystical threshold guarding an ineffable ground. It is the boundary of predicate applicability: every predicate we possess — temporal, causal, metric, and, we argue, existential — is defined on a projection of the substrate, and there is no predicate whose domain is "the beyond," because a domain is a projection and the beyond is, by construction, not one. Questions of the form "what is before, outside, or beyond the block, and what caused it" apply an axis-bound predicate to an argument outside its domain; they are ill-typed, not false, and to decline them is type-correction, which is distinct both from denial and from agnosticism. We turn the discipline on the article's own vocabulary: "completed" and "invariant" are themselves time-facet predicates and apply to the substrate only by courtesy, the substrate being strictly a-temporal rather than unchanging. We show why the beyond cannot be furnished with a simulator, a creator, or a multiverse — positing an existent beyond the block predicates existence outside its domain — and we are careful that even the apophasis avoids negation: the thesis is not that nothing lies beyond but that "beyond" is not a domain over which something and nothing range. We distinguish the position sharply from negative theology, locate Parmenides, Kant, Wittgenstein, Carnap, and the relativistic block universe as partial apprehensions of it, and mark what is here established as categorial fact and what is imported.

Keywords: categorial boundary; block universe; apophasis; predicate applicability; category error; internal and external questions; eternalism; foundations of physics; philosophy of information.

Contents

1 Introduction

A foundational theory must say where it stops, and the saying is itself a test of the theory's discipline. A theory that stops by positing a final entity — a first cause, a ground of being, a containing void — has not stopped at all but has taken one more step and called it the last. A theory that stops by professing ignorance — "beyond this we cannot know" — has left a well-formed question standing with its answer withheld, and invites the perpetual return of the question. The Information Substrate Theory stops in neither way. It stops at a boundary that is not a frontier between two regions but the edge of the applicability of its predicates, and it holds that the questions which seem to point past the boundary are not unanswered but ill-formed.

The thesis was stated in the opening article and relied upon in the second [5, 6]: within the describable, the universe is a completed, invariant structure — a block whose change lives on its time-facet without the whole being in time — and beyond the describable the theory is silent of categories, the boundary of the describable being the boundary of predicate applicability. The present article is devoted to that boundary. It asks, and answers, what kind of limit it is; it shows why it forbids furnishing the beyond with any entity; and it disciplines, with some care, the very words — "block," "completed," "invariant" — in which the thesis is naturally expressed, for those words are themselves drawn from a facet and must be applied to the whole only by courtesy.

The difficulty of the topic is that it is easy to violate the boundary in the act of describing it. To say "there is nothing beyond the block" is already to have furnished the beyond with an absence, predicating non-existence of a region one has thereby treated as a possible bearer of existence. To say "the beyond is unknowable" is to have classified it as a knowable-in-principle whose knowledge we lack. The disciplined statement is harder and stranger: that "beyond" is not a domain over which "something" and "nothing," "knowable" and "unknowable" range at all — that the locative "beyond the block" fails to denote, not because its referent is hidden but because the predicates that would fix a referent are facets of the block and have no application outside it. Making this precise, and keeping it precise even in its own statement, is the work of the article.

We proceed by fixing the block and disciplining its description (Section 3), stating the categorial boundary and the distinction between the ill-typed, the false, and the open (Section 4), showing why the beyond cannot be furnished and how this differs from negative theology (Section 5), and locating the predecessors who apprehended part of the position (Section 6). Throughout, the discipline is the one the volume requires: to say less, and to keep the saying well-formed.

2 Logical Structure of the Argument

The argument fixes an object, draws a boundary, derives a prohibition, and locates its ancestry.

Step 1 (Section 3): The block, and its self-discipline.
We fix the block as the differentiated whole available to predication, and show that the predicates by which we naturally describe it as "completed" and "invariant" are time-facet predicates that apply to it only by courtesy (Proposition 3.2). The substrate is strictly a-temporal, neither complete-at-a-time nor changing.

Step 2 (Section 4): The categorial boundary.
Every predicate we possess is defined on a projection (Proposition 4.2); the boundary of the describable is therefore the boundary of predicate applicability (Proposition 4.3). A question that applies an axis-bound predicate beyond its domain is ill-typed, which is distinct both from being false and from being open (Proposition 4.4).

Step 3 (Section 5): Why we do not furnish the beyond.
Existence-predication is internal to the block (Proposition 5.1); positing a simulator, creator, or multiverse beyond it predicates existence outside its domain and is ill-typed. The apophasis is itself stated without negation (Proposition 5.2), and is distinguished from negative theology (Remark 5.3).

Step 4 (Section 6): Locating the predecessors.
Parmenides, Kant, Wittgenstein, Carnap, and the relativistic block universe are placed as partial apprehensions of the categorial boundary, each generously and each with its over-reading marked.

Step 5 (Section 7): Epistemic status.
An explicit ledger separates what is here established as categorial fact, what is imported, and what is deferred.

The dependencies are linear: Step 2 presupposes the substrate–projection apparatus of Step 1 and of Article 1; Step 3 applies Step 2 to the existence-predicate; Steps 4–5 locate and audit. The thesis advanced by the totality is that the limit of the theory is a categorial limit, internal to the structure of predication, and that respecting it is a matter of type-correctness rather than of modesty or of reverence.

3 The Block and Its Self-Discipline

3.1 What the block is

We fix the object whose boundary is at issue.

Definition 3.1 (The describable and the block)**.** The *describable* is the totality of what can be predicated by means of the axes of the substrate: the projections $\mathcal{S}_\alpha$ and the composites built from them. The *block* is the substrate $\mathcal{S}$ together with all its projections — the differentiated whole as available to predication. Within the block, change is real on the time-facet $\mathcal{S}_{\alpha_t}$, and the physics of process proceeds there untouched.

The block is realist. Its differentiatedness does not depend on being projected, and the projection measures report facts about its facets rather than artifacts of the reporting; this realism was fixed in the opening article and is presupposed here [5]. The block is also, in the loose and natural way of speaking, *completed* and *invariant*: it is not a thing still in the making, and it does not, as a whole, change. But "completed" and "invariant" are exactly the kind of words this article must handle with care, and we handle them at once, for if we do not discipline our own description we shall have violated the boundary before reaching it.

3.2 "Completed" and "invariant" hold only by courtesy

Proposition 3.2 (Self-discipline of the descriptive vocabulary)**.** *"Completed" (not still becoming across time) and "invariant" (constant across time) are time-facet predicates: their evaluation requires the time-coordinate. By the categorial inapplicability of temporal predicates to the substrate, established in Article 1, they do not strictly apply to the block-as-substrate. They describe the block faithfully only* from within *the time-facet, where change is real and the whole that contains it does not itself change. Strictly, the substrate is* a-temporal*: neither complete-at-a-time nor changing, since both predicates require a time the substrate is not in.*

Proof. "Completed" contrasts a finished state with an earlier unfinished one and so indexes times; "invariant" asserts equality of a quantity across times; both take arguments in the time-projection $\mathcal{S}_{\alpha_t}$. The substrate is not in the time-projection but is the unreduced structure of which it is one facet [5]; applying a time-indexed predicate to it is ill-typed. Hence the words apply to the substrate not literally but by courtesy — as the best the time-bound describer can do — while strictly the substrate is a-temporal. □

This is not a quibble but the method of the article in miniature, applied first to itself. We will go on speaking of "the block," of its "completeness" and "invariance," because the time-bound describer has no better words; but the reader should hear these words throughout as courtesies extended from within the time-facet, not as literal temporal predications of a timeless whole. The same caution attaches to the spatial metaphor in "block" and to the locative in "boundary": they are facet-drawn figures for a structure that is not literally a solid in space with an edge. The figures are serviceable; their literal reading is the error the article exists to forestall.

3.3 Counterarguments and Responses

Objection 1 (if even "invariant" fails to apply, the block is being described in self-undermining terms). *Response.* It is being described in *acknowledged courtesy* terms, which is not self-undermining but candid. A time-bound describer must use time-facet words; the discipline is not to abandon them — there are no others — but to flag their status, so that they are not pressed into literal claims about the substrate. A description that marks its own figurative reach is more rigorous than one that does not, not less.

Objection 2 (this reduces to the trivial point that the substrate is outside time, already made in Article 1). *Response.* It extends that point in a way the article needs: it shows that the *vocabulary of the boundary itself* — "completed," "invariant," "block," "beyond" — is facet-drawn, and therefore that stating the boundary correctly requires constant courtesy-marking. Article 1 established the timelessness of the substrate; the present proposition turns that result on the act of description, which is what makes a disciplined statement of the boundary possible at all.

3.4 Section Result and Implications

We have fixed the block as the differentiated whole available to predication (Definition 3.1) and shown that the predicates by which we describe it as completed and invariant apply to it only by courtesy, the substrate being strictly a-temporal (Proposition 3.2). The implication is that the description of the boundary must everywhere mark the figurative status of its facet-drawn vocabulary. With the block fixed and our own words disciplined, we can state the boundary.

4 The Categorial Boundary

4.1 Predicates are projection-internal

Definition 4.1 (Axis-bound predicate)**.** A predicate P is *axis-bound* if its domain is an axis-projection $\mathcal{S}_\alpha$ or a composite of such projections: P is meaningfully applied only to arguments that are elements or tuples of some $\mathcal{S}_\alpha$.

Proposition 4.2 (Every predicate we possess is projection-internal)**.** *The predicates available to a describer — temporal (before, after, during), causal (produces, depends on), metric (near, far, how much), and the existential (exists, obtains), treated in Section 5 — are axis-bound: each is defined on a projection of the substrate. None is defined on a region "beyond" the block, because a domain of predication is a projection, and "beyond the block" is, by construction, not a projection of the block.*

Proof. By survey of the kinds. Temporal predicates take arguments in the time-projection (Article 1); metric predicates take arguments in the relevant geometric projection; causal predicates relate states across the time-projection and so are likewise time-indexed; the existential predicate is shown projection-internal in Proposition 5.1. In each case the domain is a projection $\mathcal{S}_\alpha$. A putative argument "beyond the block" is, by Definition 3.1, not an element of any $\mathcal{S}_\alpha$, since the projections exhaust the describable; hence no axis-bound predicate has it in its domain, and there is no further, non-axis-bound predicate available to a describer whose every predicate is drawn from a facet. □

The proposition does not claim that the list of axes is closed or that no further predicate could ever be introduced; it claims that every predicate we *have* is projection-internal, and that any predicate introduced in the future, being a predicate of a describer, would be defined on some projection and so would inherit the same limitation. The boundary is thus not a contingent gap in our current vocabulary but a structural feature of predication as such, on the substrate account of it.

4.2 The boundary of the describable is the boundary of predicate applicability

Proposition 4.3 (The categorial boundary)**.** *A question of the form "what is P beyond / before / outside the block?", where P is any axis-bound predicate, applies P to an argument outside its domain. Such a question is ill-typed: it does not express a proposition with a truth-value. The boundary of the describable coincides with the boundary of predicate applicability.*

Proof. By Proposition 4.2, P is defined only on a projection $\mathcal{S}_\alpha$, and "beyond the block" denotes no element of any $\mathcal{S}_\alpha$. The application $P(\text{beyond})$ is therefore the application of a function to an argument outside its domain, which is undefined; the resulting string is not a false proposition but a non-proposition. Since this holds for every axis-bound predicate, and all our predicates are axis-bound, no well-formed proposition concerns the beyond, and the describable — the range of well-formed predication — terminates exactly where predicate applicability does. □

4.3 Ill-typed, false, and open are three different things

The single most important distinction in this article is the one that keeps the categorial boundary from collapsing into either denial or agnosticism.

Proposition 4.4 (Ill-typed ≠ false ≠ open)**.** *An ill-typed question has no truth-value. It is therefore not* false*: falsity is a property of well-formed propositions, and an ill-typed string expresses none. It is also not* open*: an open question has a determinate truth-value that is merely unknown, whereas an ill-typed question has no truth-value to be known or unknown. To decline an ill-typed question is* type-correction*, a third response distinct from both asserting its negation and suspending judgment on it.*

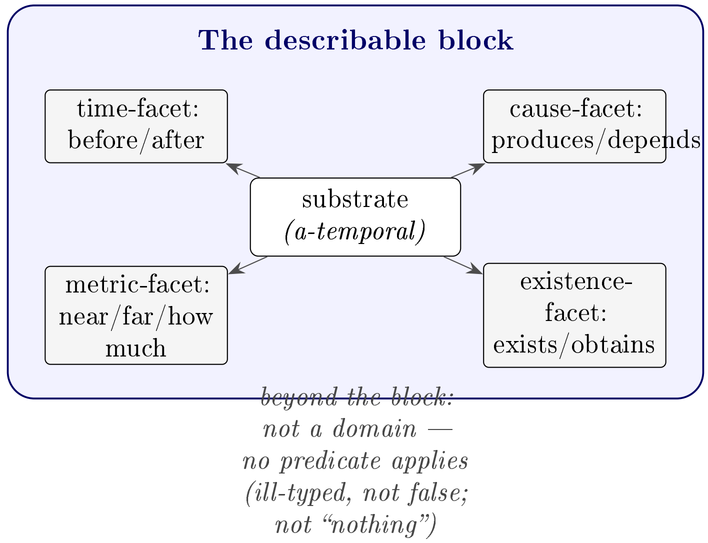

Figure 1: The boundary as the edge of predicate applicability. Within the block, every predicate — temporal, causal, metric, existential — is defined on a projection (facet) of the a-temporal substrate. The rectangle's edge is not a wall in space with a region behind it but the limit beyond which no predicate has a domain. "Beyond the block" fails to denote: questions that apply a facet-predicate there are ill-typed, not false, and the disciplined statement is not that nothing lies beyond but that "beyond" is not a domain over which something and nothing range (Proposition 5.2).

Proof. Falsity, truth, and openness presuppose a proposition: to be false is to be a proposition that fails to hold; to be open is to be a proposition whose holding is undetermined for the inquirer. An ill-typed string, by Proposition 4.3, is not a proposition, so none of these statuses applies. The appropriate response is neither assertion of a negation (which would be a further ill-typed string, "P does not hold beyond the block") nor suspension of judgment (which presupposes a proposition to suspend judgment on) but the diagnosis that the string fails to express a proposition — type-correction. □

The distinction does real work against two opposite misreadings. Against the reading that IST *denies* a beyond — asserts that nothing lies beyond the block — the proposition shows that the denial would itself be ill-typed, predicating non-existence of a non-domain; IST asserts no such thing. Against the reading that IST is merely *agnostic* — that there may be a beyond about which we cannot know — it shows that agnosticism presupposes a well-formed question whose answer is withheld, which is precisely what the beyond does not supply. IST is neither denying nor withholding; it is correcting the type of the question.

4.4 Counterarguments and Responses

Objection 1 (this is verificationism in new dress: "what cannot be checked is meaningless"). *Response.* It is not verificationism, and the difference is principled.

Verificationism ties meaning to a method of confirmation and founders on its own statement and on well-formed-but-unconfirmable claims. The present boundary is not epistemic but categorial: a question past the boundary is ill-typed not because we lack a method to settle it but because its predicate has no domain there. A claim can be perfectly unverifiable and yet well-formed, if its predicate has a domain; and a claim can be ill-typed independently of any method of checking. Meaning here turns on predicate domains, not on verification.

Objection 2 (declaring questions ill-typed is a convenient way to dodge hard metaphysics). *Response.* It would be, if the ill-typing were asserted rather than shown. The article does not stipulate that boundary-questions are malformed; it derives their malformedness from the projection-internal character of predicates (Proposition 4.2), which is in turn grounded in the substrate apparatus of Article 1. A reader who would resist the conclusion must contest that apparatus, not merely accuse the conclusion of convenience. The diagnosis is a result, with premises that can be attacked, not a refusal.

4.5 Section Result and Implications

We have established that every predicate we possess is projection-internal (Proposition 4.2), that the boundary of the describable is therefore the boundary of predicate applicability (Proposition 4.3, Figure 1), and that questions past the boundary are ill-typed — a status distinct from both falsity and openness (Proposition 4.4). The implication, drawn out next, is that the beyond cannot be furnished with any entity, because to furnish it would be to predicate existence across the very boundary that existence-predication does not cross.

5 Why the Beyond Cannot Be Furnished

5.1 Existence-predication is block-internal

The temptation a foundational theory must resist is to take one more step past its boundary and install an entity there: a first cause, a creator, a containing multiverse, a simulator running the world. The categorial boundary forbids this, and not as a matter of taste or parsimony but as a matter of type.

Proposition 5.1 (Existence-predication is internal to the block)**.** *Within the describable, for a thing to* exist *is for it to be a differentiated part of the block — to figure among the configurations or projected entities of $\mathcal{S}$. The predicate "exists" is thus defined on the block; it is axis-bound in the sense of Definition 4.1, its domain being the projections that constitute the describable. There is no further, framework-external sense of "exists" available to a describer. Consequently, to assert that some entity exists* beyond *the block is to apply "exists" to an argument outside its domain, which is* *ill-typed.*

Proof. On the substrate account, to be is to be differentiated (Article 1); to exist, for a particular, is to be a differentiated part of the whole, registered along some axis as a configuration or a projected entity. The predicate "exists" therefore has the describable as its domain. A would-be entity "beyond the block" is, by Definition 3.1, not a differentiated part of the block, hence not in the domain of "exists"; predicating existence of it applies the predicate outside its domain and is ill-typed by Proposition 4.3. No second sense of "exists," detached from the block, is available, for any sense a describer commands is drawn from the projections that constitute the describable. □

This is the precise reason the beyond cannot be furnished. A simulator, a creator, a multiverse beyond the block are all *existents posited beyond the block*, and the positing predicates existence across a boundary that existence-predication does not cross. The prohibition is therefore not the modest "we have no evidence for such entities" nor the parsimonious "we should not multiply them"; it is the categorial "the assertion that they exist is ill-typed." One cannot furnish a region that is not a domain.

5.2 Apophasis without negation

The boundary must be stated without committing, in the statement, the very error it diagnoses. This requires a discipline finer than ordinary denial.

Proposition 5.2 (The apophasis is stated without negation)**.** *The disciplined statement of the boundary is not "nothing exists beyond the block" but "the question of existence beyond the block is ill-typed." The former predicates non-existence of the beyond and so treats "beyond" as a domain over which existence and its absence range, recommitting the error of Proposition 5.1. The latter makes no predication of the beyond at all; it classifies a string as failing to express a proposition.*

Proof. "Nothing exists beyond the block" has the logical form $\neg\exists x\,(x \text{ is beyond the block})$, which quantifies over a domain "beyond the block" and denies it has members; it thereby treats the beyond as a domain, contrary to Proposition 4.3. The disciplined statement instead ascends one level: it does not quantify over the beyond but evaluates the *type* of the string "x is beyond the block," finding it ill-formed. Ascending from object-level predication to the typing of the string avoids any predication of the beyond, which is the only way to state the boundary without crossing it. □

The point is delicate and worth dwelling on, because it is where careless apophasis fails. To say "there is nothing beyond" feels like the maximally modest claim, but it is in fact a substantive and ill-typed one: it pictures a beyond and reports it empty. The genuinely modest — and the only well-formed — move is to decline to picture the beyond at all, by noting that the words purporting to picture it do not, on inspection, express a proposition. The boundary is held not by asserting the emptiness of the beyond but by declining the locative "beyond" as non-denoting.

5.3 Distinction from negative theology, and the relocation of the simulator

Remark 5.3 (Not negative theology)**.** Negative theology approaches a real ground — the divine, the Absolute — by the denial of predicates, holding that the ground exceeds all of them while nonetheless being, and being the source of, the world. The present position is its categorial opposite. It does not approach a ground that exceeds predication; it denies that "beyond the block" names a domain at all, and so denies that there is a ground-entity behind the block to be approached, whether by negation or otherwise. The apophasis here is anti-metaphysical, not reverent: it removes the beyond as a subject of discourse rather than honoring it as an ineffable one. Where negative theology says "the ground is, but no predicate captures it," IST says "there is no well-formed question to which a ground would be the answer."

Remark 5.4 (Why the simulator must live in the model)**.** The categorial boundary settles, in advance, the form any defensible version of the simulation hypothesis must take. A simulator posited *outside* the block is an existent beyond the block and is ill-typed by Proposition 5.1; the beyond cannot house it. If the simulation hypothesis is to be well-formed at all, the simulation must be located *within* the block, in a model-building subsystem — which is exactly the relocation that the later articles of this volume carry out, placing the simulation in the model and not in the world [7]. The boundary does not refute the simulation hypothesis; it constrains where a coherent version of it may live, and the only available place is inside the block.

5.4 Counterarguments and Responses

Objection 1 (surely it is coherent to wonder what caused the universe; physics asks this). *Response.* Physics asks, and answers, causal questions *within* the time-projection: it relates earlier states to later ones, including very early ones. That inquiry is well-formed and is untouched here. What is ill-typed is the extrapolation of the causal predicate to "the cause of the block as a whole," which would relate the whole to something outside the time-projection it constitutes — a causal predication with no domain for its second argument. The boundary forbids the extrapolation, not the physics; cosmology within the time-facet is as legitimate as ever.

Objection 2 (declining to furnish the beyond is indistinguishable in practice from asserting the block is all there is, which is a substantive metaphysical claim). *Response.* The two are distinguishable exactly where it matters. "The block is all there is" quantifies over a totality and excludes a remainder, predicating non-existence of the remainder — ill-typed by Proposition 5.2. Declining to furnish the beyond makes no such quantification; it neither affirms nor denies a remainder but classifies the locative "beyond" as non-denoting. In practice this means IST will not assert "and there is nothing else," any more than it asserts "and there is something else"; it asserts neither, because both are ill-typed.

5.5 Section Result and Implications

We have shown that existence-predication is internal to the block (Proposition 5.1), so that furnishing the beyond with any entity — simulator, creator, multiverse — is ill-typed rather than merely unwarranted; that the boundary must be stated without negation, on pain of re-committing the error (Proposition 5.2); and that the position is the categorial opposite of negative theology (Remark 5.3), with the consequence that a coherent simulation hypothesis must locate its simulation inside the block (Remark 5.4). The implication is that the theory's refusal to overreach is enforced by its own type-discipline, not by a policy of restraint. We turn to the predecessors who saw part of this.

6 Locating the Predecessors

The categorial boundary has a long ancestry, and the stance of this volume is to locate each ancestor as a partial apprehension rather than to claim novelty against them. Several saw the boundary clearly in one respect and over-read it in another; the locating notes both.

Kirk et al. [4] report the argument of *Parmenides* that what-is is one, ungenerated, and complete, and that coming-to-be and perishing are names without application to it. This is an early and remarkable apprehension of the a-temporal whole — of Proposition 3.2's point that becoming does not apply to the whole. Parmenides over-read it into the denial of plurality and change altogether; IST keeps change and plurality where they live, on the facets, and reserves the a-temporality for the whole. The insight is the timelessness of being-as-a-whole; the over-reading is the erasure of the facets in which time and multiplicity are real.

Kant [3] diagnosed, in the antinomies of pure reason, that questions such as whether the world has a beginning in time or a limit in space generate contradiction because they extend the categories — causality, magnitude — beyond the conditions of their legitimate use. This is the nearest classical statement of the present boundary: a category misapplied beyond its domain yields not a hard truth but a malformed question. IST's debt to Kant here is direct. The difference is that Kant grounds the limit in the conditions of possible experience, within a transcendental idealism; IST grounds it in the axis–projection structure of differentiatedness, and the block is realist, not a phenomenon constituted by the knower. The diagnosis of category-misapplication is shared; its foundation is relocated from the knower to the structure.

Wittgenstein [9] drew the limit at language: the limits of language are the limits of the world, and whereof one cannot speak, thereof one must be silent. This is the apophatic insight in its purest modern form, and IST shares its conclusion that the boundary is a limit of sense rather than a frontier between regions. IST specifies what the *Tractatus* leaves gnomic: the silence is categorial, owed to the projection-internal character of predicates, and its discipline is type-correction rather than mystical gesture. Where

Wittgenstein's silence can read as reverence before the unsayable, IST's is the flat observation that the unsayable strings are not propositions.

Carnap [1] is the closest methodological ancestor. His distinction between questions *internal* to a linguistic framework, which have answers, and *external* questions about the framework as a whole, which are pseudo-questions or at most practical, is almost exactly the present distinction between projection-internal predication and the ill-typed locative "beyond the block." IST refines Carnap in one respect that matters: where Carnap grounds the internal/external line in a conventional choice of framework, so that the line could have been drawn elsewhere, IST grounds it in the substrate–projection structure, so that the line is not a convention but a categorial fact about predication. The internal/external distinction is inherited; its status is upgraded from convention to structure.

Finally, the block universe of relativity — the completed four-dimensional spacetime of Minkowski [8], and the eternalism Gödel [2] drew from it — is a genuine apprehension of the block, and a physical one. On the present reading it is the block seen along the time-facet taken whole: spacetime as a completed four-dimensional structure is the time-projection rendered total, and it is correctly completed and invariant in the courtesy-sense of Proposition 3.2. IST locates it as such and adds only that this four-dimensional block is itself a projection — the time-facet writ large — of the axis-free substrate, of which it is one rendering among the facets. The relativistic block is the deepest *physical* apprehension of the boundary; the substrate block is the structure of which even it is a facet.

6.1 Counterarguments and Responses

Objection 1 (claiming Kant and Carnap as ancestors while rejecting their foundations is to take the conclusion and discard the argument). *Response.* It is to take a diagnosis and supply it a different foundation, which is a normal and legitimate move when the diagnosis is sound but its original grounding is contestable. Kant's diagnosis of category-misapplication and Carnap's internal/external distinction survive detachment from transcendental idealism and from framework-conventionalism respectively; IST supplies them the substrate–projection grounding and claims only that grounding as its own, crediting the diagnoses to their authors.

Objection 2 (the relativistic block is a substantive physical theory, not a "facet," and subsuming it understates it). *Response.* The subsumption understates nothing. The four-dimensional block is affirmed in full as the correct physical rendering of the time-facet taken whole; calling it the time-facet writ large is not a demotion but a placement. IST adds a claim — that this block is a projection of an axis-free substrate — which is contentful and defended elsewhere in the volume (Article 3); it subtracts nothing from the physics of spacetime.

6.2 Section Result and Implications

We have located Parmenides, Kant, Wittgenstein, Carnap, and the relativistic block universe as partial apprehensions of the categorial boundary — the a-temporal whole, the misapplied category, the limit of sense, the internal/external distinction, and the completed four-dimensional structure — crediting each insight and marking each over-reading or relocation. The implication is that the boundary is not an idiosyncrasy of this volume but a recurring discovery, here given a foundation in the substrate and a discipline in type-correctness. We close with the ledger.

7 Epistemic Status: Established, Imported, Deferred

In the manner of the series, we close with an explicit ledger. *Established (here)* marks the categorial results proved in this article; *Established (imported)* marks premises and results taken from prior articles and not re-argued; *Located* marks the placement of a predecessor; *Deferred* marks a claim handled in a named later article.

Claim	**Status**	**Locus**
Temporal predicates inapplicable to substrate	Established (imported)	Kriger [5]
Substrate–projection apparatus; realism of the block	Established (imported)	Kriger [5]
"Completed," "invariant" hold only by courtesy	Established (here)	Prop. 3.2
Every predicate we possess is projection-internal	Established (here)	Prop. 4.2
Boundary of describable = boundary of predicate applicability	Established (here)	Prop. 4.3
Ill-typed $\neq$ false $\neq$ open	Established (here)	Prop. 4.4
Existence-predication is block-internal	Established (here)	Prop. 5.1
The beyond cannot be furnished (ill-typed, not unwarranted)	Established (here)	Prop. 5.1
Apophasis stated without negation	Established (here)	Prop. 5.2
Position is not negative theology	Established (here)	Rem. 5.3
Predecessors (Parmenides, Kant, Wittgenstein, Carnap, block)	Located	§6
A coherent simulation lives in the model, not the beyond	Deferred	Article 13
Known measures descend from the substrate	Imported / deferred	Article 3

The ledger marks the article's discipline. It *proves* categorial facts — that predicates are projection-internal, that the boundary is the boundary of their applicability, that the beyond cannot be furnished, and that the apophasis must be stated without negation. It *imports* the timelessness and realism of the substrate from Article 1. It *locates* its ancestors. And it *defers* to the later articles the positive relocation of the simulation into the model, which the boundary makes not optional but required.

8 Discussion

The result of this article is a clarification of the theory's own limit. The boundary at which the Information Substrate Theory stops is not a frontier with a region behind it, not a horizon of present ignorance, and not a mystical threshold; it is the boundary of predicate applicability, fixed by the projection-internal character of every predicate a describer commands. Questions that seem to point past it — what is before, outside, or beyond the block, and what caused it — are ill-typed, a status the article is at pains to distinguish from both falsity and openness. The beyond cannot be furnished with a simulator, a creator, or a multiverse, because to furnish it would be to predicate existence across the boundary that existence-predication does not cross; and even the statement of the boundary must avoid negation, since "nothing lies beyond" would itself treat the beyond as a domain.

It is worth marking the relation between this boundary and the impossibility of nothing on which the volume's descent began [5]. The two are different and must not be conflated. The impossibility of nothing concerns the null state — the absence of all being — and rules it out as a pseudo-alternative to there being something; it is a modal claim about a candidate state. The categorial boundary concerns the locative "beyond the block" and finds it non-denoting; it is a claim about the type of a string. The null is a would-be state shown impossible; the beyond is a would-be region shown to be no region. A reader who runs them together will hear the boundary as the claim that "there is nothing beyond," which is exactly the ill-typed negation Proposition 5.2 forbids. The discipline of the volume requires holding the two apart: being cannot fail to be differentiated, and "beyond the differentiated whole" does not denote.

The boundary is, finally, what keeps the substrate honest. A theory that grounds the world in a substrate is under permanent temptation to let the substrate swell into a metaphysical entity behind the world — a ground in the theological sense, a thing on which the block rests. The categorial boundary forecloses the swelling: the substrate is the axis-free differentiatedness of the describable block, and there is, in the strict sense the article has labored to make precise, no place for it to rest and nothing for it to rest on, because "place" and "rest on" are facet-predicates with no application beyond the block. The substrate is not a deeper entity; it is the whole, taken without any axis, and the boundary is the guarantee that it stays that and does not become a god by another name.

8.1 Limitations

Three limitations are material. *First*, the categorial results rest on the projection-internal character of predicates (Proposition 4.2), which is grounded in the substrate apparatus of Article 1; a reader who rejects that apparatus may grant the internal/external *form* of the boundary, after Carnap, while denying its upgrade from convention to categorial fact. *Second*, the claim that *every* predicate a describer commands is projection-internal is defended by survey and by the structural observation that any new predicate would be a describer's predicate and so projection-bound; it is not a closed-form theorem, and a predicate genuinely defined on no projection, were one exhibited, would bound the result. *Third*, the article makes no positive claim whatever about the beyond, and is therefore, by design, unable to satisfy a reader who wants the boundary either crossed or declared empty; that dissatisfaction is the intended effect of type-correction and not a defect to be remedied.

8.2 Future Directions

The article's sequels are internal to the volume and to its programme. The positive relocation of the simulation hypothesis into a model-building subsystem, which the boundary makes mandatory rather than optional, is carried out in Part III [7]; the present article supplies the categorial reason the relocation is forced. Beyond the volume, the open question the boundary bequeaths is whether the projection-internal character of predicates can be established other than by survey — whether there is a structural theorem to the effect that any predicate definable by a subsystem of the block has a projection for its domain. Such a theorem would convert Proposition 4.2 from a well-supported survey into a result, and with it would harden the categorial boundary from a disciplined stance into a demonstrated feature of predication on the substrate. That question is stated here and left to the work that can address it.

Acknowledgments

The author thanks colleagues of the Information Physics Institute and the Institute of Integrative and Interdisciplinary Research for sustained discussion of the foundations of the series.

References

[1] R. Carnap. Empiricism, semantics, and ontology. *Revue Internationale de Philosophie*, 4(11):20–40, 1950.

[2] K. Gödel. A remark about the relationship between relativity theory and idealistic philosophy. In P. A. Schilpp, editor, *Albert Einstein: Philosopher-Scientist*, pages 555–562. Open Court, La Salle, IL, 1949.

[3] I. Kant. *Critique of Pure Reason.* Cambridge University Press, Cambridge, 1998. Translated and edited by P. Guyer and A. W. Wood; original 1781/1787.

[4] G. S. Kirk, J. E. Raven, and M. Schofield. *The Presocratic Philosophers.* Cambridge University Press, Cambridge, 2nd edition, 1983.

[5] B. Kriger. The information substrate: differentiation as the ground of information. The Information Substrate Theory (IST), Volume VII, Article 1. Zenodo, 2026.

[6] B. Kriger. Substrate and projection: why the universe needs no information. The Information Substrate Theory (IST), Volume VII, Article 2. Zenodo, 2026.

[7] B. Kriger. Simulation in the model, not in the world: relocating the simulation hypothesis. The Information Substrate Theory (IST), Volume VII, Article 13. Zenodo, 2026.

[8] H. Minkowski. Raum und Zeit. *Physikalische Zeitschrift*, 10:75–88, 1909.

[9] L. Wittgenstein. *Tractatus Logico-Philosophicus.* Routledge & Kegan Paul, London, 1961. Translated by D. F. Pears and B. F. McGuinness; original 1921.

A A Typology of Malformed Boundary-Questions

The diagnoses of this article can be collected as a typology of the questions that appear to point past the boundary, with the predicate each misapplies and the type-correction each receives. The table is offered not to dismiss the questions — they are natural and recur — but to make their category explicit, so that the difference between an ill-typed question and an open one is visible at a glance.

Question	**Predicate misapplied**	**Type-correction**
"What was there *before* the block?"	temporal (before)	"Before" indexes the time-facet; the block is not in it (Prop. 3.2). Ill-typed, not open.
"What is *outside* the block?"	metric / locative (outside)	"Outside" indexes a spatial projection; "beyond the block" is no element of any projection (Prop. 4.3). Non-denoting.
"What *caused* the block?"	causal (produces)	Causation relates states across the time-facet; the whole has no second argument outside it (Prop. 4.2). Ill-typed.
"Does anything *exist* beyond it?"	existential (exists)	"Exists" is block-internal; predicating it beyond the block applies it outside its domain (Prop. 5.1). Ill-typed.
"Is there *nothing* beyond it?"	existential (negated)	Denying existence beyond still treats "beyond" as a domain (Prop. 5.2). Equally ill-typed; the apophasis takes no such form.
"What *simulates* the block?"	existential + causal	A simulator beyond the block is an existent posited beyond it (Prop. 5.1); a coherent simulation lives in a subsystem (Rem. 5.4).

The typology exhibits the uniform structure of the malformed questions: each takes a predicate at home on some facet — before, outside, caused, exists — and applies it to the whole or to a would-be region beyond the whole, where the predicate has no domain. The correction is in each case the same: not an answer, and not a confession of ignorance, but the observation that the question fails to express a proposition. The last two rows are the ones most easily mistaken: "is there nothing beyond?" looks like the safe, modest question, but it is as ill-typed as its affirmative twin, and "what simulates the block?" is well-formed only when its simulator is sought inside the block, never beyond it.

The Axis of the Channel: Shannon Information as a Substrate Projection

The complete theory of one facet, located: entropy, mutual information, and capacity as constructs on the channel reduction

Boris Kriger[1,2]

[1] Information Physics Institute
boris.kriger@informationphysicsinstitute.net

[2] Institute of Integrative and Interdisciplinary Research, Toronto
boriskriger@interdisciplinary-institute.org

ORCID: https://orcid.org/0009-0001-0034-2903

The Information Substrate Theory (IST) — Volume VII, Article 5
Part II: The Projections — classification of information theories

Abstract

This article opens the classificatory part of the volume by locating Shannon's theory of information on the channel axis of the substrate. The Projection Theorem of Article 3 already derived Shannon entropy as the channel-axis projection, fixed up to the choice of unit by the Shannon–Khinchin characterization; the present article develops that placement in full and shows that the entire apparatus of classical information theory — joint and conditional entropy, mutual information, relative entropy, channel capacity, and the source and channel coding theorems — consists of constructs on the channel reduction. The channel reduction retains the statistics of repetition of distinguished outcomes and discards their identity and the individual realized on a given draw; Shannon's measure is the registered differentiation per draw of the source so reduced. We exhibit the shape of the axis: its native indeterminacy is the unit, its composition law is additivity, and the Rényi and Tsallis families are deformed channel axes that collapse to Shannon at their additive point, their relation to Shannon being a relation among channel reductions in the sense of Article 3. We then dissolve a long debate: the celebrated observer-dependence of Shannon entropy is not a subjectivity of the functional but the dependence of the channel *reduction* on a prior choice of which outcomes count as distinct; given the reduction, the entropy is objective, and the spectator is spent in selecting the reduction, not in evaluating the measure. The data-processing inequality is identified as the channel-axis face of the lossiness of reduction. Finally we mark what the channel axis cannot see — the relations to the description and macrostate axes that Article 3 supplies — and locate Shannon's theory generously as the complete and correct theory of one facet, not a rival to be demoted but a projection to be placed.

Keywords: Shannon entropy; channel axis; mutual information; relative entropy; channel capacity; Rényi and Tsallis entropy; maximum entropy; data-processing inequality; foundations of physics.

Contents

1 Introduction

Of all the measures that carry the name of information, Shannon's is the one whose theory is most complete. Where Kolmogorov complexity is uncomputable and effective complexity definitionally contested, Shannon's entropy comes with a settled axiomatic characterization, a pair of operational coding theorems, and a rich apparatus — conditional entropy, mutual information, relative entropy, channel capacity — that has organized communication engineering for three quarters of a century [1, 7]. A classificatory account of the information measures that did not do justice to the completeness of Shannon's theory would be worthless. The aim of this article is to do that justice while making good on the central claim of the volume: that Shannon's theory, complete as it is, is the complete theory of *one facet* — the channel axis — and that locating it as such subtracts nothing from it and adds the relations it cannot, from within, see.

The placement itself is already established. The Projection Theorem of Article 3 [4] derived Shannon entropy as the channel-axis projection of the substrate: the channel reduction sends the differentiated source to a probability vector, and the Shannon–Khinchin characterization then fixes the only continuous, additive, refinement-respecting measure of registered differentiation on that vector, up to the choice of unit. The present article does not re-prove that derivation; it takes it as the point of departure and develops the channel axis in the detail a dedicated treatment permits. We ask what the channel reduction retains and discards, exhibit the full Shannon apparatus as constructs on the reduced object, describe the shape of the axis — its unit, its composition law, and the deformed families that surround Shannon's additive point — and then turn to two matters on which the substrate reading makes a genuine contribution beyond the derivation.

The first is a debate as old as the theory: whether Shannon entropy is objective or observer-relative. The objectivist points to the determinacy of H given a distribution; the subjectivist, following the maximum-entropy tradition [2], points to the ineliminable role of the modeler's choice of variables and constraints. The substrate reading resolves the dispute by locating the observer precisely: the observer-dependence is the dependence of the channel *reduction* on a prior choice of which outcomes count as distinct and how the source is modeled, and given that reduction the functional H adds no further observer and is objective. The spectator is spent in choosing the reduction, not in evaluating the measure. The second contribution is to identify the data-processing inequality — that processing cannot increase mutual information — as the channel-axis face of the lossiness of reduction, connecting Shannon's theory to the asymmetric bridge that the later articles of the volume develop.

We close by marking, without complaint, what the channel axis cannot see. Shannon's theory is complete within its axis and silent about the relations of the channel projection to the others — to the description axis of Kolmogorov complexity, to the macrostate axis of thermodynamic entropy. Those relations are exactly the rigidities the Projection Theorem supplies, and they are invisible to a theory that takes the channel measure as primitive. This is the sense in which the substrate frame holds Shannon's theory rather than rivals it: it assigns the theory its facet and explains why the facet relates to the others as it does.

2 Logical Structure of the Argument

The argument places the channel reduction, builds the apparatus on it, describes the shape of the axis, makes two contributions, and marks the limit.

Step 1 (Section 3): The channel reduction.
We make precise what the channel reduction retains — the statistics of repetition of distinguished outcomes — and what it discards — the identity of outcomes and the individual realized on a draw. The source-as-ensemble is the channel projection.

Step 2 (Section 4): The channel-axis apparatus.
Shannon entropy is recalled as the channel-axis projection, unique up to unit (Article 3); we then exhibit joint and conditional entropy, mutual information (Proposition 4.2), relative entropy (Proposition 4.3), capacity, and the coding theorems as further constructs on the channel reduction.

Step 3 (Section 5): The shape of the axis.
The native indeterminacy is the unit and the composition law is additivity; the Rényi and Tsallis families are deformed channel axes collapsing to Shannon at the additive point (Proposition 5.1, Figure 1), their relation to Shannon being induced in the sense of Article 3.

Step 4 (Section 6): Reduction-relativity.
The observer-dependence of H is the dependence of the channel reduction on a prior choice (Proposition 6.1); given the reduction, H is objective. The data-processing inequality is the channel-axis face of the lossiness of reduction (Proposition 6.2).

Step 5 (Section 7): What the channel axis cannot see.
Shannon's theory is complete within its axis and blind to the cross-axis relations the Projection Theorem supplies; we locate the theory generously as the complete theory of one facet.

Step 6 (Section 8): Epistemic status.
An explicit ledger separates what is proved here, what is imported, and what is affirmed as standard information theory.

The dependencies are linear: Steps 2–4 operate on the reduction of Step 1; Step 5 sets the whole within the cross-axis frame of Article 3. The thesis advanced is that Shannon's theory is the complete and correct theory of the channel facet, and that seeing it so locates its observer-relativity, its deformed relatives, and its limits without diminishing it.

3 The Channel Reduction

3.1 What the channel reduction retains and discards

We begin from the reduction that defines the axis, stated in Article 3 and developed here.

Definition 3.1 (Channel reduction)**.** The *channel reduction* ρ_{ch} sends the differentiated source $\mathcal{S}$ to a finite probability vector $p = (p_1, \ldots, p_m)$ on its m distinguished outcomes, retaining how often each outcome is distinguished from the rest under repetition and discarding everything else. The channel projection $\mathcal{S}_{\mathrm{ch}}$ is the simplex of finite probability vectors.

Two things are discarded by ρ_{ch}, and naming them fixes the character of the axis. The first is the *identity* of the outcomes: the channel projection records that there are m distinguished outcomes with weights $p_1, \ldots, p_m$, but not which configuration is which; relabeling the outcomes leaves the channel projection unchanged, which is why Shannon entropy is a symmetric function of the p_i. The second is the *individual* realized on a given draw: the channel projection is a property of the source as a repeatable ensemble, not of any single emission. These two discardings distinguish the channel axis sharply from the description axis of Kolmogorov complexity, which retains exactly the individual that the channel axis discards (Article 3); the two axes are reductions of the same substrate along orthogonal coordinates, one toward the statistics of repetition, the other toward the specification of an individual.

3.2 The source as a repeatable ensemble

The channel reduction presupposes that the source is the kind of thing that repeats — that "how often" has application. This is not a trivial presupposition, and it is the substantive content of choosing the channel axis. To take the channel reduction is to treat the differentiated source as an ensemble: a population of distinguishable outcomes whose relative frequencies, in the limit of repetition, are the weights p_i. Where the source genuinely repeats — a transmitter emitting symbols, a physical system sampled many times — the channel reduction is natural and the frequencies are empirical. Where it does not, the weights are supplied by a model, and the modeler's role, examined in Section 6, is exactly the supply of the ensemble the channel reduction requires. Either way, the channel projection is the source-as-ensemble, and the measure built on it measures a property of the ensemble, never of the substrate behind it (Article 2).

3.3 Counterarguments and Responses

Objection 1 (the channel reduction simply is a probability space; the substrate talk is idle). *Response.* The probability vector appears only after ρ_{ch}; it is the output of the reduction, not the substrate. The point of locating it as a reduction is precisely that a different reduction of the same source — the description reduction, the macrostate reduction — yields an object with no probabilities in it, or different ones, and that the relations among these reductions are the content of Article 3. The substrate is what is common to the reductions and is poorer than any of them; the channel probability space is one of its images.

Objection 2 (discarding outcome identity is a defect, since real channels care which symbol is which). *Response.* Real communication systems care about identity, and they recover it with structure the channel reduction does not by itself supply — a codebook, a semantics. That recovery is a further construction, and where it matters it is added explicitly (as in the joint reductions of Section 4). The symmetry of H under relabeling is not a defect but an exact record of what the bare channel reduction retains: the statistics of repetition, not the labels.

3.4 Section Result and Implications

We have fixed the channel reduction as the map to the source-as-ensemble, retaining the statistics of repetition and discarding outcome identity and the individual draw (Definition 3.1). The implication is that every quantity built on the channel projection is a property of an ensemble, symmetric under relabeling and indifferent to the individual — a profile that the Shannon apparatus exhibits exactly, as we now show.

4 The Channel-Axis Apparatus

4.1 The functional, recalled

The functional native to the channel axis is fixed, up to unit, by characterization. We recall the result of Article 3 without re-proving it.

Proposition 4.1 (Shannon as channel-axis projection; 4)**.** *The unique family of functionals on finite probability vectors that is continuous, maximal at the uniform vector, expansible, and additive under refinement (the Shannon–Khinchin conditions) is*

$$H(p) = -\lambda \sum_{i=1}^{m} p_i \log p_i, \qquad \lambda > 0, \tag{1}$$

and Shannon information is $H \circ \rho_{\mathrm{ch}}$, the channel-axis projection of the substrate, with λ (equivalently the logarithm base) the native unit of the axis.

The interpretation, established in Article 3 and presupposed here, is that H measures the registered differentiation per draw of the source-as-ensemble: it is large when the ensemble is richly distinguished (many outcomes of comparable weight) and small when it is nearly settled (one outcome dominant). Everything that follows is built on this functional and on the joint and conditional reductions of two or more sources.

4.2 Joint, conditional, and mutual constructs

Two channel reductions of a joint source $\mathcal{S}$ — one to a variable X, one to a variable Y — yield a joint probability vector $p(x, y)$ and its marginals. The standard constructs follow.

Proposition 4.2 (Mutual information as cross-reduction shared differentiation)**.** *For two channel sub-reductions X, Y of a joint source, the mutual information*

$$I(X;Y) = H(X) + H(Y) - H(X,Y) = \sum_{x,y} p(x,y) \log \frac{p(x,y)}{p(x)p(y)} \tag{2}$$

satisfies $I(X;Y) \geq 0$, with equality iff the two sub-reductions are independent. It is the registered differentiation shared between the two channel projections.

Proof. Non-negativity is the non-negativity of relative entropy (Proposition 4.3) applied to $p(x, y)$ against $p(x)p(y)$; equality holds iff $p(x, y) = p(x)p(y)$, i.e. independence. The identity with $H(X) + H(Y) - H(X, Y)$ is algebraic [1]. □

On the substrate reading, mutual information is a quantity *within* the channel axis that relates two of its sub-reductions: it measures how much the differentiation registered by X overlaps that registered by Y. It is not a cross-*axis* relation in the sense of Article 3 — both X and Y are channel reductions — but it is the channel axis's own way of comparing two of its projections, and it is the quantity on which the operational theory of communication is built.

4.3 Relative entropy as reduction-mismatch

Proposition 4.3 (Relative entropy as the cost of the wrong reduction-reference)**.** *For probability vectors p, q on the same outcomes,*

$$D_{\mathrm{KL}}(p \,\|\, q) = \sum_i p_i \log \frac{p_i}{q_i} \geq 0, \tag{3}$$

with equality iff $p = q$ [5]. It measures the excess channel-axis cost — in bits per draw — of describing the source p as though its ensemble were q.

Proof. Non-negativity is Gibbs' inequality, a convexity consequence of log [1]; equality holds iff the distributions coincide. The operational reading as excess code length follows from the source coding theorem (Section 4.4): coding for q when the source is p costs $H(p) + D_{\mathrm{KL}}(p\|q)$ bits per symbol instead of $H(p)$. □

Relative entropy is the channel axis's measure of a *mismatch between reductions*: q is a candidate channel reduction (a model of the ensemble) and $D_{\mathrm{KL}}(p\|q)$ is the penalty, in the axis's own units, for using it when the true ensemble is p. This makes it the natural channel-axis instrument for the modeler's situation of Section 6, where the reduction is supplied rather than read off.

4.4 Capacity and the operational theorems

The operational heart of Shannon's theory is two theorems, which we state as the content that makes the channel axis a theory of communication and not merely of a functional.

The *source coding theorem* states that the source-as-ensemble can be encoded, with vanishing error, at an average rate approaching H bits per symbol and no lower; this is the operational meaning of the entropy as the registered differentiation per draw. The *channel coding theorem* states that information can be transmitted reliably over a noisy channel at any rate below the *capacity*

$$C = \max_{p(x)} I(X;Y), \tag{4}$$

the maximum mutual information over input distributions, and not above it [1, 7]. Both theorems are statements about what the channel reduction permits: the first about how far the source-as-ensemble can be compressed, the second about how much shared differentiation can be forced through a noisy reduction. On the substrate reading they are the operational characterization of the channel axis, and they make no reference to, and need none of, the other axes.

4.5 Counterarguments and Responses

Objection 1 (mutual information is surely a cross-axis relation, since it relates two variables). *Response.* It relates two *channel* reductions, X and Y, both of which live on the channel axis; it is a comparison internal to the axis, not a relation between the channel axis and another axis such as the description or macrostate axis. The cross-axis relations of Article 3 relate H to K or to S; mutual information relates $H(X)$ to $H(Y)$. The distinction matters because the rigidity theorems of Article 3 govern the former, while mutual information is governed by the ordinary identities of the channel axis.

Objection 2 (presenting the coding theorems without proof understates them). *Response.* The coding theorems are landmarks and we do not slight them; we state them as the operational content of the channel axis and refer to the standard treatments for proof [1], because the article's task is classification, not re-derivation of established information theory. Their role here is to show that the channel axis carries not only a functional but a complete operational theory, which is exactly what makes Shannon's the most developed of the projections.

4.6 Section Result and Implications

We have exhibited the Shannon apparatus — entropy, joint and conditional entropy, mutual information (Proposition 4.2), relative entropy (Proposition 4.3), capacity, and the coding theorems — as constructs on the channel reduction, all internal to the channel axis and all indifferent to the other axes. The implication is that the channel axis is not a bare functional but a complete operational theory of one facet. We now describe the shape of that axis: its unit, its composition law, and the family of deformed axes that surround it.

5 The Shape of the Axis

5.1 Unit and composition law

Two structural features of the channel axis were used in Article 3 and are made explicit here. The *native indeterminacy* of the axis is the unit λ of Equation (1), equivalently the base of the logarithm: bits, nats, or any other choice. No fact about the channel projection fixes the unit, because the reduction discards the scale along with the labels; the unit is supplied externally, and its choice is the channel-axis analogue of the choice of universal machine on the description axis (Article 3). The *composition law* of the axis is additivity: for independent sources, $H(p \otimes q) = H(p) + H(q)$, which is the content of the grouping axiom in the characterization and the property that singles out Shannon's functional among its deformed relatives. These two features — a free unit and an additive composition — are what the rescaling rigidity of Article 3 used to force the proportionality of thermodynamic to channel entropy, $S = k_B H$, with k_B the unit ratio.

5.2 The deformed channel axes

Shannon's functional is the additive member of a family. Relaxing the additive composition law while keeping continuity, maximality, and expansibility yields one-parameter families of channel-type measures, and locating them is part of placing the channel axis.

Proposition 5.1 (Deformed channel axes collapse to Shannon at the additive point)**.** *The Rényi entropy [6]*

$$H_\alpha(p) = \frac{1}{1-\alpha} \log \sum_i p_i^\alpha, \qquad \alpha \geq 0,\ \alpha \neq 1, \tag{5}$$

and the Tsallis entropy [8] $S_q(p) = \frac{1}{q-1}\big(1 - \sum_i p_i^q\big)$ each define a one-parameter family of channel-type measures with a deformed composition law, and each recovers Shannon's H in the additive limit: $H_\alpha \to H$ as $\alpha \to 1$ and $S_q \to H$ (in nats) as $q \to 1$. The Rényi family is non-increasing in α, with $H_0 = \log m$ the Hartley (maximum) entropy and $H_\infty = -\log \max_i p_i$ the min-entropy.

Proof. The limits are standard: applying l'Hôpital's rule to (5) as $\alpha \to 1$ gives $-\sum_i p_i \log p_i$, and likewise for S_q as $q \to 1$ [6, 8]. Monotonicity of H_α in α and

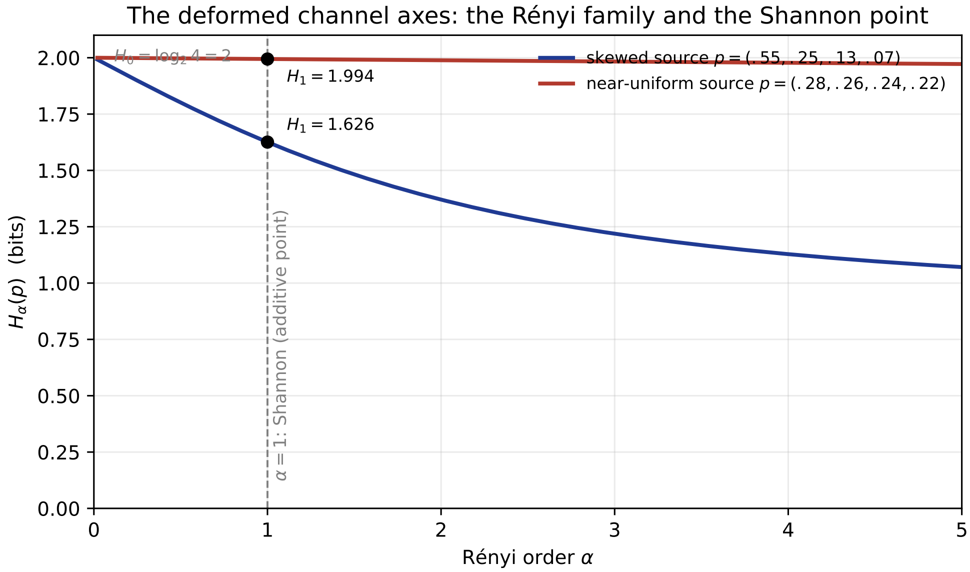

Figure 1: The deformed channel axes. The Rényi family $H_\alpha(p)$ for two sources, as a function of the order α. Shannon entropy sits at the additive point $\alpha = 1$ (black dots), where the deformed composition law becomes additive; the Hartley maximum $H_0 = \log_2 4 = 2$ bits is shared by all four-outcome sources, and the family decreases monotonically toward the min-entropy as $\alpha \to \infty$. The one-parameter family is a family of channel reductions whose composition law is deformed away from additivity, all collapsing to Shannon's reduction at $\alpha = 1$; their relation to Shannon is a relation among channel reductions, induced in the sense of Article 3.

the endpoint values H_0, H_∞ are elementary properties of (5). The deformed composition laws — additive for Rényi of independent sources, q-additive for Tsallis — replace the Shannon grouping axiom, which is what removes these families from the characterization of Proposition 4.1 and leaves Shannon as the additive member. $\square$

On the substrate reading, the deformed families are not rivals to Shannon's measure but *neighboring channel axes*: each takes the same channel reduction — the source-as-ensemble — but measures its registered differentiation under a deformed composition law. Article 3's clause (ii) anticipated exactly this: the convergence of the Rényi family to Shannon as $\alpha \to 1$ is a relation among channel reductions, the one-parameter family collapsing to the additive point, and so is induced rather than coincidental. Where a deformed measure is the natural one — as the Tsallis entropy is for certain long-range-correlated systems — the substrate frame locates it as the channel axis under a non-additive composition, neither privileging nor demoting it relative to Shannon's additive member.

5.3 Counterarguments and Responses

Objection 1 (if the deformed families are equally good channel measures, the characterization of Shannon is arbitrary). *Response.* It is not arbitrary; it is conditional on the composition law. Shannon's functional is the unique channel measure that is *additive* under independent composition; the deformed families are the unique channel measures under their respective deformed laws. The characterization selects Shannon once additivity is fixed, and the choice of additivity is the choice of the additive channel axis among the family. Article 3's rescaling rigidity shows why the additive axis is the one that relates by proportionality to the macrostate axis; the deformed axes relate differently, which is itself informative.

Objection 2 (Tsallis entropy is controversial in physics; endorsing it overreaches). *Response.* We endorse nothing about the physical domains in which a deformed entropy is appropriate; that is an empirical question outside this article. We claim only the structural fact that the deformed families are channel-type measures with non-additive composition laws collapsing to Shannon at the additive point. Where they are physically apt and where they are not is for the physics of the system to decide; the substrate frame locates them structurally without adjudicating their domains.

5.4 Section Result and Implications

We have exhibited the shape of the channel axis: its native indeterminacy is the unit, its composition law is additivity, and the Rényi and Tsallis families are deformed channel axes collapsing to Shannon's additive member (Proposition 5.1, Figure 1), with their convergence to Shannon induced in the sense of Article 3. The implication is that Shannon's measure occupies a distinguished — the additive — point on a family of channel measures, and that its distinction is structural, not arbitrary. We turn to the role of the observer.

6 Reduction-Relativity

6.1 Observer-relativity is reduction-relativity

The longest-running dispute about Shannon entropy is whether it is an objective property of a system or a subjective measure of an observer's ignorance. The substrate reading dissolves the dispute by locating the observer exactly.

Proposition 6.1 (The observer-dependence of H is the reduction-dependence of ρ_{ch})**.** *The value of $H \circ \rho_{\mathrm{ch}}$ depends on the observer only through the channel reduction ρ_{ch}: the prior choice of which configurations count as the distinguished outcomes and what weights the source-as-ensemble carries. Given ρ_{ch} — given the partition and the distribution — the functional H is determined and contains no further observer. The* *spectator is spent in selecting the reduction, not in evaluating the measure.*

Proof. $H \circ \rho_{\mathrm{ch}}(\mathcal{S})$ is the composition of a fixed functional H (Proposition 4.1) with the reduction ρ_{ch}. Any dependence on the observer must enter through one of the two factors. The functional H is fixed up to the unit by characterization and admits no observer parameter; hence all observer-dependence enters through ρ_{ch}, which encodes the choice of distinguished outcomes and their weights. Fixing ρ_{ch} fixes p, and $H(p)$ is then determined. □

This is the resolution. The objectivist is right that H is objective: given the reduction, the entropy is a determinate number, not a matter of opinion. The subjectivist is right that a choice is ineliminable: the reduction must be selected, and different reductions of the same source yield different entropies. Both are right because they are speaking of different factors of the same composition. The maximum-entropy tradition [2] is located precisely: choosing the distribution of maximum entropy subject to constraints is choosing the least committal channel reduction compatible with what one has fixed, and Jaynes's principle is a principle for *selecting the reduction*, not for evaluating the functional. The apparent subjectivity of entropy is the necessity of choosing a reduction; the objectivity is the determinacy of the functional once the reduction is chosen.

6.2 The data-processing inequality as the lossiness of reduction

A second feature of the channel axis connects it to the asymmetric bridge of the volume.

Proposition 6.2 (Data-processing inequality as reduction-lossiness)**.** *If $X \to Y \to Z$ is a Markov chain — Z obtained from Y by further processing, independent of X given Y — then $I(X;Z) \leq I(X;Y)$ [1]. Processing cannot increase the shared differentiation; a further reduction can only lose, never create, mutual information.*

Proof. Standard: the chain rule for mutual information gives $I(X;Y,Z) = I(X;Z) + I(X;Y \mid Z) = I(X;Y)+I(X;Z \mid Y)$, and the Markov condition makes $I(X;Z \mid Y) = 0$; non-negativity of $I(X;Y \mid Z)$ then yields $I(X;Z) \leq I(X;Y)$ [1]. □

On the substrate reading, the data-processing inequality is the channel-axis face of a theme that runs through the whole volume: reduction loses. A processing step $Y \to Z$ is a further reduction of the channel projection, and the inequality states that such a reduction cannot increase the differentiation Y shares with X — it can only preserve or discard it. This is the channel-axis instance of the asymmetry that Part III of the volume develops into the arrow of time and the incompleteness of measurement: physical projection loses, and the data-processing inequality is the precise, proved statement of the loss within the channel axis. The model that would recover what processing discarded must, as the bridge articles argue, add structure by hypothesis; the channel axis by itself only records that the loss cannot be undone by further processing.

6.3 Counterarguments and Responses

Objection 1 (locating the observer in the reduction just relabels the subjectivity, it does not remove it). *Response.* It does not claim to remove the choice; it claims to locate it, and the location is the resolution of the dispute, not its evasion. The dispute was whether H is objective or subjective, and the answer is that it is objective given the reduction and that the reduction requires a choice — a precise division that neither disputant had drawn. Relabeling that makes a long-standing dispute dissolve into two compatible truths is explanatory, not cosmetic.

Objection 2 (the data-processing inequality is elementary; calling it the "lossiness of reduction" dresses it up). *Response.* The inequality is elementary, and we prove it in two lines. What the substrate reading adds is not difficulty but *connection*: it identifies the inequality as the channel-axis instance of the lossy-projection asymmetry that organizes the whole volume, so that the arrow of time (Part III) and the data-processing inequality are seen as the same asymmetry on different axes. The value is the unification, not a complication of the proof.

6.4 Section Result and Implications

We have located the observer-dependence of Shannon entropy as the reduction-dependence of the channel projection, dissolving the subjective/objective dispute (Proposition 6.1), and identified the data-processing inequality as the channel-axis face of the lossiness of reduction (Proposition 6.2). The implication is that the two features of Shannon's theory most often read as philosophical embarrassments — its observer-relativity and the irreversibility of processing — are, on the substrate reading, exact structural facts about the channel reduction and its further reductions. We close by marking what the axis cannot see.

7 What the Channel Axis Cannot See

The channel axis is a complete operational theory of one facet, and its completeness is real: within the axis, entropy, mutual information, capacity, and the coding theorems answer the questions the axis can pose. But a theory that takes the channel measure as primitive is, by that very choice, blind to the relations between the channel projection and the others.

It cannot see, from within, why the expected shortest description length of a source equals its entropy — the agreement of the channel axis with the description axis — because the description axis does not appear in it; that agreement is the averaging rigidity of Article 3, visible only from the substrate. It cannot see why thermodynamic entropy is k_B times an information measure — the proportionality of the channel axis to the macrostate axis — because the macrostate axis does not appear in it; that proportionality is the rescaling rigidity of Article 3. It records these agreements, where it meets them, as fortunate coincidences or as bridges to other theories; it cannot derive them, because the common domain that induces them is exactly what a single-axis theory omits.

This is the precise sense in which the substrate frame holds Shannon's theory rather than competes with it. The frame assigns the theory its facet — the channel axis, with its ensemble reduction, its additive functional, its operational theorems — and affirms its completeness there. It then supplies what the facet cannot: the relations to the other facets, as the rigidities of the morphisms between reductions. Shannon's theory is not a rival to be demoted but a projection to be placed, and placing it is the whole of the present article. We locate Shannon, with the generosity the volume owes its predecessors, as the author of the complete and correct theory of the channel facet — the most developed of the projections, and the one whose very completeness made the substrate behind it easy to overlook.

7.1 Counterarguments and Responses

Objection 1 (if Shannon's theory is complete within its axis, the substrate adds nothing a working scientist needs). *Response.* It adds nothing *within* the channel axis, and that is the point of calling the axis complete. It adds what a working scientist needs precisely when crossing axes: when relating a compression bound to an entropy, or an erasure cost to a bit count, or a thermodynamic entropy to a channel entropy. Those cross-axis relations are where single-axis theories report coincidences, and they are where the substrate frame supplies derivations (Article 3). The addition is not within the channel axis but at its borders with the others.

Objection 2 (calling Shannon's completeness a reason the substrate was "overlooked" is rhetorical). *Response.* It is a historical observation with a structural basis. A facet whose own theory is complete and self-contained gives no internal prompt to look beneath it; the very success of the channel axis as a closed operational theory is why the question "what is it a projection of?" did not press itself. The observation explains, without disparaging anyone, why the substrate reading is a late rather than an early move.

7.2 Section Result and Implications

We have marked what the channel axis cannot see — the cross-axis relations to the description and macrostate axes that Article 3 supplies — and located Shannon's theory as the complete and correct theory of the channel facet, held by the substrate frame rather than rivaled by it. The implication, carried to the rest of Part II, is that each projection is to be treated as Shannon's has been here: as a complete theory of its facet, located on its axis, with its relations to the other facets reserved for the substrate. The next articles do for the description, regularity, and macrostate axes what this one has done for the channel.

8 Epistemic Status: Proved, Imported, Affirmed

In the manner of the series, we close with an explicit ledger. *Proved (here)* marks results established in this article; *Imported* marks results taken from prior articles, chiefly the derivation and rigidity theorems of Article 3; *Affirmed (standard)* marks established information theory the article endorses and locates without re-deriving.

Claim	**Status**	**Locus**
Shannon = channel-axis projection, unique up to unit	Imported	Kriger [4] (Art. 3)
Rescaling / averaging rigidities (cross-axis relations)	Imported	Kriger [4] (Art. 3)
Channel reduction retains repetition, discards identity/individual	Proved (here)	Def. 3.1
Mutual information as cross-reduction shared differentiation	Proved (here)	Prop. 4.2
Relative entropy as reduction-mismatch cost	Proved (here)	Prop. 4.3
Deformed families collapse to Shannon at additive point	Proved (here)	Prop. 5.1
Observer-dependence is reduction-dependence	Proved (here)	Prop. 6.1
Data-processing inequality as reduction-lossiness	Proved (here)	Prop. 6.2
Entropy / capacity / coding theorems	Affirmed (standard)	Cover and Thomas [1], Shannon [7]
Maximum-entropy principle (located as reduction-selection)	Affirmed (standard)	Jaynes [2]
Rényi and Tsallis entropies	Affirmed (standard)	Rényi [6], Tsallis [8]
Cross-axis relations invisible to a single-axis theory	Proved (here)	§7

The ledger marks the division of labor. The article *imports* the derivation and the cross-axis rigidities from Article 3; it *proves* the structural facts about the channel reduction, the apparatus built on it, the deformed families, and the two reduction-facing results; and it *affirms*, without re-deriving, the established information theory it locates.

9 Discussion

The result of this article is a placement, not a new theorem of information theory. Shannon's theory is located on the channel axis as the complete operational theory of one facet: the channel reduction sends the differentiated source to a repeatable ensemble, the Shannon functional measures the registered differentiation per draw, and the apparatus of mutual information, relative entropy, capacity, and the coding theorems are constructs on that reduction. The placement does two things a working information theorist may value beyond the bookkeeping. It dissolves the subjective/objective dispute about entropy by locating the observer in the reduction and the objectivity in the functional, showing the disputants to have been speaking of different factors of one composition. And it identifies the data-processing inequality as the channel-axis instance of the lossy-projection asymmetry that the volume develops elsewhere into the arrow of time, so that an elementary inequality and a cosmological arrow are seen as one asymmetry on different axes.

The placement also continues the discipline of the volume. It does not demote Shannon's theory; it affirms its completeness within its axis and reserves only the cross-axis relations for the substrate. It does not promote the channel measure to the ground, as the reification pictures of Article 2 would; the bit is located as a channel-axis quantity, and the source-as-ensemble as a projection, never as the substrate. And it marks honestly the limit of the axis: that a theory which takes the channel measure as primitive cannot see why its measure agrees, where it does, with the measures of the other axes, because the common domain that induces the agreement is exactly what a single-axis theory omits.

9.1 Limitations

Three limitations are material. *First*, the article imports the derivation of Shannon as the channel-axis projection and the cross-axis rigidities from Article 3 and does not re-establish them; a reader who doubts those results will find the placement only as secure as its source. *Second*, the treatment of the deformed families is structural, locating Rényi and Tsallis entropies as channel axes under deformed composition laws without adjudicating the physical domains in which a deformed measure is the appropriate one, which is an empirical matter outside the article's scope. *Third*, the operational theorems are stated and located but not proved, the article's task being classification rather than the re-derivation of established information theory; their landmark status is acknowledged and referred to the standard treatments.

9.2 Future Directions

The article is the first of Part II, and its sequels do for the other projections what it has done for the channel axis: the description axis of Kolmogorov complexity, the regularity axis of effective complexity, and the macrostate axis of thermodynamic entropy each receive the same located treatment in the articles that follow. Beyond the volume, the placement of the deformed families invites a precise question continuous with Article 3: whether the convergence of the Rényi and Tsallis families to Shannon, here described as a relation among channel reductions collapsing to the additive point, is induced by a morphism of the kind the Projection Theorem classifies — a one-parameter deformation morphism — and whether such deformation morphisms belong in the typology of inter-axis maps alongside rescaling, kernel, and refinement. That question connects the present located treatment to the open completeness conjecture of Article 3 and is left to the work that can address it.

Acknowledgments

The author thanks colleagues of the Information Physics Institute and the Institute of Integrative and Interdisciplinary Research for sustained discussion of the foundations of the series.

References

[1] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Wiley-Interscience, Hoboken, NJ, 2nd edition, 2006.

[2] E. T. Jaynes. Information theory and statistical mechanics. *Physical Review*, 106(4):620–630, 1957.

[3] B. Kriger. The information substrate: differentiation as the ground of information. The Information Substrate Theory (IST), Volume VII, Article 1. Zenodo, 2026.

[4] B. Kriger. The projection theorem: how known information measures descend from the substrate. The Information Substrate Theory (IST), Volume VII, Article 3. Zenodo, 2026.

[5] S. Kullback and R. A. Leibler. On information and sufficiency. *Annals of Mathematical Statistics*, 22(1):79–86, 1951.

[6] A. Rényi. On measures of entropy and information. In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability*, volume 1, pages 547–561. University of California Press, 1961.

[7] C. E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423, 1948.

[8] C. Tsallis. Possible generalization of Boltzmann-Gibbs statistics. *Journal of Statistical Physics*, 52(1-2):479–487, 1988.

A Code for Figure 1

The following Python script generates Figure 1, computing the Rényi entropy family $H_\alpha(p) = \frac{1}{1-\alpha} \log_2 \sum_i p_i^\alpha$ for two four-outcome sources and marking the Shannon value at the additive point $\alpha = 1$.

```
import numpy as np

def shannon(p):
    p = np.asarray(p, float); p = p[p > 0]
    return -np.sum(p * np.log2(p))

def renyi(p, alpha):
    p = np.asarray(p, float); p = p[p > 0]
    if abs(alpha - 1.0) < 1e-9:
        return shannon(p)
    return (1.0 / (1.0 - alpha)) * np.log2(np.sum(p ** alpha))

p_skew = np.array([0.55, 0.25, 0.13, 0.07])   # Shannon H_1 = 1.6256 bits
p_near = np.array([0.28, 0.26, 0.24, 0.22])   # Shannon H_1 = 1.9942 bits

alphas = np.concatenate([np.linspace(0.02, 0.98, 60),
                         np.linspace(1.02, 5.0, 80)])
H_skew = [renyi(p_skew, a) for a in alphas]   # H_0 = 2 (Hartley), decreasing
H_near = [renyi(p_near, a) for a in alphas]   # toward the min-entropy as a->inf
```

The Axis of Description: Kolmogorov Complexity and the Algorithmic Projection

The complete theory of the individual object, located: incompressibility, the universal distribution, and uncomputability

Boris Kriger[1,2]

[1] Information Physics Institute
boris.kriger@informationphysicsinstitute.net

[2] Institute of Integrative and Interdisciplinary Research, Toronto
boriskriger@interdisciplinary-institute.org

ORCID: https://orcid.org/0009-0001-0034-2903

The Information Substrate Theory (IST) — Volume VII, Article 6
Part II: The Projections — classification of information theories

Abstract

Where Shannon's theory is the complete theory of the channel facet, Kolmogorov complexity is the complete theory of the description facet: the measure not of a source as a repeatable ensemble but of an individual object as something to be specified. This article, the second of the classificatory part of the volume, locates the algorithmic theory of information on the description axis of the substrate. The Projection Theorem of Article 3 already derived Kolmogorov complexity as the description-axis projection, fixed up to an additive constant by the invariance theorem; here we develop that placement in full. The description reduction retains the full individual and discards the ensemble, exactly the orthogonal complement of the channel reduction, and the algorithmic apparatus — plain and prefix complexity, conditional and joint complexity, the chain rule, algorithmic mutual information with its symmetry, the universal distribution, and Solomonoff induction — consists of constructs on the retained individual. We exhibit the shape of the axis: its native indeterminacy is the additive machine constant, and its composition law is subadditivity up to a logarithmic term rather than the exact additivity of the channel axis, a difference that explains why the channel–description relation is a kernel (averaging) morphism and not a rescaling. Levin's coding theorem, $-\log m(x) = K(x) + O(1)$, is identified as the per-object face of that averaging relation, and randomness is located as incompressibility: the individual the description axis cannot compress. We then treat the characteristic limitation of the axis — uncomputability — as the price of retaining the whole individual: the shortest description of a retained individual is upper semicomputable but not computable, a structural feature bearing on evaluation, not on the derivation of the measure. Finally we mark what the description axis cannot see and locate Kolmogorov, Solomonoff, Chaitin, and Levin as the authors of the complete theory of one facet.

Keywords: Kolmogorov complexity; algorithmic information; description axis; prefix complexity; universal distribution; algorithmic randomness; Solomonoff induction; uncomputability; foundations of physics.

Contents

1 Introduction

The companion article on the channel axis located Shannon's theory as the complete theory of one facet: the source taken as a repeatable ensemble, with entropy the registered differentiation per draw [6]. There is a second theory of information, as deep and as complete, that measures something the channel theory discards entirely. Kolmogorov complexity is not a property of a source but of an individual object: the length of the shortest program that, run on a fixed universal machine, outputs that object and halts [2, 3, 10]. Where Shannon's entropy answers "how much differentiation does this source register per draw," Kolmogorov's complexity answers "how much specification does this individual require." The two questions are about orthogonal things — an ensemble and an individual — and the present article locates the algorithmic theory on the description axis of the substrate, as the complete and correct theory of the second.

The placement is established. The Projection Theorem of Article 3 [5] derived Kolmogorov complexity as the description-axis projection: the description reduction sends the differentiated substrate to a single configuration regarded as a finite object to be reconstructed, and the invariance theorem then fixes the measure — the shortest description — up to an additive constant depending only on the choice of universal machine. As with the channel axis, this article does not re-prove the derivation but develops the axis in the detail a dedicated treatment permits. We ask what the description reduction retains and discards, exhibit the full algorithmic apparatus as constructs on the retained individual, describe the shape of the axis, and then treat the feature that most sharply distinguishes the description axis from the channel axis: not observer-relativity but uncomputability.

Two contributions go beyond the bare derivation, paralleling the two that the channel article made. The first concerns the *composition law*. The channel axis is exactly additive — the entropy of independent sources adds — and Article 3 used that exact additivity to force the proportionality of thermodynamic to channel entropy by a rescaling. The description axis is additive only up to a logarithmic term: the complexity of a pair is the sum of the complexities plus an $O(\log)$ correction, never exactly. We show that this inexactness is not a blemish but the structural reason the channel–description relation is a *kernel* (averaging) morphism rather than a rescaling: the description axis does not enter a proportionality with the channel axis, it enters an expectation, and Levin's coding theorem is the per-object statement of that expectation. The second contribution concerns uncomputability. We locate it as the price of the axis's defining choice: the description reduction retains the whole individual, and the shortest specification of a whole individual is not effectively findable. This is the description axis's characteristic limitation, the structural counterpart of the channel axis's observer-relativity, and like it a feature to be located rather than a defect to be apologized for.

We close, as the channel article did, by marking what the description axis cannot see — the relations to the channel and regularity axes that the substrate supplies — and by locating the theory's authors generously. Kolmogorov, Solomonoff, Chaitin, and Levin built the complete theory of the individual object; the substrate frame assigns it its facet and explains why that facet relates to the others as it does.

2 Logical Structure of the Argument

The argument parallels that of the channel article, axis for axis.

Step 1 (Section 3): The description reduction.
We make precise what the description reduction retains — the full individual — and discards — the ensemble — and show it to be the orthogonal complement of the channel reduction.

Step 2 (Section 4): The algorithmic apparatus.
Kolmogorov complexity is recalled as the description-axis projection, unique up to an additive constant (Article 3); we then exhibit prefix complexity, conditional and joint complexity and the chain rule, algorithmic mutual information and its symmetry (Proposition 4.2), the universal distribution and Levin's coding theorem (Proposition 4.3), and Solomonoff induction.

Step 3 (Section 5): The shape of the axis.
The native indeterminacy is the additive machine constant; the composition law is subadditivity up to a logarithmic term (Proposition 5.1), which explains why the channel–description relation is a kernel morphism, not a rescaling. Randomness is located as incompressibility (Proposition 5.2, Figure 1).

Step 4 (Section 6): Uncomputability.
K is upper semi-computable but not computable (Proposition 6.1); the shortest description of a retained individual is not effectively findable. This bears on evaluation, not on the derivation of the measure, and is read as the inverse of a lossy projection.

Step 5 (Section 7): What the description axis cannot see.
The axis is complete within itself and blind to its relations to the channel axis (the averaging rigidity) and the regularity axis (the refinement decomposition); we locate the theory's authors generously and forward the regularity decomposition to Article 7.

Step 6 (Section 8): Epistemic status.
An explicit ledger separates what is proved here, what is imported, and what is affirmed as standard algorithmic information theory.

The dependencies are linear: Steps 2–4 operate on the reduction of Step 1; Step 5 sets the whole within the cross-axis frame of Article 3. The thesis advanced is that Kolmogorov complexity is the complete and correct theory of the description facet, and that seeing it so locates its composition law, its randomness, and its uncomputability without diminishing it.

3 The Description Reduction

3.1 What the description reduction retains and discards

We begin from the reduction that defines the axis, stated in Article 3 and developed here.

Definition 3.1 (Description reduction)**.** The *description reduction* ρ_{de} sends the differentiated substrate $\mathcal{S}$ to a single configuration x, regarded as a finite object to be reconstructed: a binary string, retaining the entire individual and discarding everything about the source from which it was drawn. The description projection $\mathcal{S}_{\mathrm{de}}$ is the set of finite objects, and the measure on it is the length of the shortest program that outputs x.

Two things are discarded by ρ_{de}, and they are exactly the two things the channel reduction retained. The first is the *ensemble*: the description projection records the individual x in full but says nothing about how often x, or any alternative, would occur under repetition; there is no probability vector in $\mathcal{S}_{\mathrm{de}}$. The second is *repetition* itself: the description axis is a property of one object, not of a population, and "how often" has no application to it. Conversely, the description reduction retains precisely what the channel reduction discarded — the identity of the individual, in full detail, down to the last bit. The two axes are reductions of the same substrate along orthogonal coordinates: the channel axis toward the statistics of repetition with the individual erased, the description axis toward the individual with the statistics erased.

Remark 3.2 (Orthogonality of the two axes)**.** The orthogonality is not a metaphor but a statement about what each reduction is a function of. The channel projection is invariant under relabeling the outcomes and depends only on the multiset of weights; the description projection is invariant under nothing — every bit of the individual matters — and depends on no weights at all. A quantity that is a function of one projection is, in general, not a function of the other, and the relation between them is therefore not an identity but a morphism, the averaging morphism of Article 3, developed in Section 5.

3.2 The object as specificand

To take the description reduction is to treat the differentiated object as a *specificand*: something whose reconstruction from a shortest instruction is the quantity of interest. This presupposes a fixed means of reconstruction — a universal machine — just as the channel reduction presupposed a fixed sense in which the source repeats. The machine is to the description axis what the ensemble model is to the channel axis: the structure the reduction requires, supplied externally, and contributing to the measure only an additive constant (Section 5). The measure built on the description projection is a property of the individual relative to that machine, never of the substrate behind it

(Article 2): the substrate is not a long string with a shortest program; the string is one of its projections.

3.3 Counterarguments and Responses

Objection 1 (the description reduction is just "encode x as a string"; the substrate talk is idle). *Response.* The string appears only after ρ_{de}, and the point of locating it as a reduction is that the same substrate admits a channel reduction with no individual in it and a macrostate reduction with neither the full individual nor the source. The relations among these reductions are the content of Article 3; the substrate is what is common to them and poorer than any. The algorithmic measure is one of its images, not its intrinsic size.

Objection 2 (retaining "the whole individual" is what makes the measure uncomputable, so the reduction is ill-chosen). *Response.* Retaining the whole individual is exactly the choice that defines the description axis, and its uncomputability (Section 6) is the price of that choice, not a sign that the choice is wrong. A theory of the individual object that refused to retain the individual would not be a theory of the description facet at all. The uncomputability is located as the characteristic limitation of the axis, on a par with the observer-relativity of the channel axis, and is not a reason to abandon the reduction.

3.4 Section Result and Implications

We have fixed the description reduction as the map to the individual-as-specificand, retaining the whole individual and discarding the ensemble and repetition (Definition 3.1), and shown it to be the orthogonal complement of the channel reduction (Remark 3.2). The implication is that every quantity built on the description projection is a property of an individual, sensitive to every bit and indifferent to any source — a profile the algorithmic apparatus exhibits exactly, as we now show.

4 The Algorithmic Apparatus

4.1 The functional, recalled

The functional native to the description axis is fixed, up to an additive constant, by the invariance theorem. We recall the result of Article 3 without re-proving it.

Proposition 4.1 (Kolmogorov complexity as description-axis projection; 5)**.** *For a fixed universal machine U, the complexity $K_U(x)$ is the length of the shortest program p with $U(p) = x$. For any two universal machines U, V there is a constant $c_{U,V}$, independent of x, with $|K_U(x) - K_V(x)| \leq c_{U,V}$. Kolmogorov complexity is $K \circ \rho_{\mathrm{de}}$, the* *description-axis projection of the substrate, with the machine constant the native indeterminacy of the axis.*

Throughout we use the *prefix* complexity $K(x)$, defined with respect to a universal machine whose set of valid programs is prefix-free (no program is a prefix of another). The prefix convention [2, 7] is preferred because it makes the programs satisfy the Kraft inequality, $\sum_x 2^{-K(x)} \leq 1$, so that $2^{-K(x)}$ behaves like a sub-probability and the apparatus of the next subsections — the chain rule, the universal distribution, the coding theorem — holds cleanly. The plain complexity $C(x)$ differs from $K(x)$ by at most a logarithmic term and shares the invariance property; we use K unless the distinction matters.

4.2 Conditional, joint, and the chain rule

Relative to a fixed machine, one defines the conditional complexity $K(x \mid y)$ — the length of the shortest program that outputs x given y — and the joint complexity $K(x, y)$ of a pair. These satisfy the chain rule up to a logarithmic term:

$$K(x, y) = K(x) + K(y \mid x^*) + O(1) = K(x) + K(y \mid x) + O(\log K(x, y)), \tag{1}$$

where x^* is a shortest program for x [8]. The logarithmic term is the first appearance of the inexactness that distinguishes the description axis from the exactly additive channel axis (Section 5); it is the cost of describing where one object ends and the next begins, which the prefix convention controls but does not eliminate.

4.3 Algorithmic mutual information and its symmetry

The description axis has its own analogue of Shannon's mutual information.

Proposition 4.2 (Symmetry of algorithmic information)**.** *Define the algorithmic mutual information $I_A(x : y) = K(x) - K(x \mid y^*) = K(x) + K(y) - K(x, y) + O(1)$. Then*

$$I_A(x : y) = I_A(y : x) + O(\log K(x, y)), \tag{2}$$

the shared algorithmic information of two individuals, symmetric up to a logarithmic term [8].

Proof. Substituting the chain rule (1) into $K(x) - K(x \mid y^*)$ and into $K(y) - K(y \mid x^*)$ gives both equal to $K(x) + K(y) - K(x, y) + O(1)$, whence the two conditional forms agree up to the logarithmic terms carried by the chain rule. The result is due to Gács and to Levin [8]. □

On the substrate reading, the symmetry of algorithmic information is the description axis's own way of comparing two of its objects: how much of the specification of x is already contained in y, and vice versa, the two being equal up to the logarithmic slack that the inexact additivity of the axis imposes. It is, like Shannon's mutual information, a quantity internal to its axis — a comparison of two description projections, not a cross-axis relation — and the logarithmic slack is the description-axis signature, absent from the exact identities of the channel axis.

4.4 The universal distribution and the coding theorem

The deepest construct on the description axis, and the one that links it to the channel axis, is the universal distribution.

Proposition 4.3 (Levin's coding theorem). *Let* $m(x) = \sum_{p:\, U(p)=x} 2^{-|p|}$ *be the algorithmic probability of* x *— the probability that the universal prefix machine outputs* x *on random program bits. Then*

$$-\log m(x) = K(x) + O(1) : \tag{3}$$

the negative logarithm of the algorithmic probability equals the prefix complexity up to a constant [7, 8].

Proof. That $-\log m(x) \leq K(x)$ is immediate, since the single shortest program contributes $2^{-K(x)}$ to the sum defining $m(x)$. The reverse inequality, $K(x) \leq -\log m(x) + O(1)$, is the substantive direction: it constructs a prefix code for x of length $-\log m(x) + O(1)$ from the semimeasure m, using the Kraft inequality guaranteed by the prefix convention [7, 8]. □

The coding theorem is the per-object face of the relation between the description axis and the channel axis. It says that the shortest description of an individual and the algorithmic probability of that individual are the same quantity, up to a constant; and since the algorithmic probability is a genuine (semi-)distribution, averaging K over a computable source returns the source's Shannon entropy up to a constant — the averaging rigidity $\mathbb{E}[K] = H + O(1)$ of Article 3, exact at the rate level by Brudno's theorem [1, 12]. We develop this link in Section 5; here it suffices to mark that the universal distribution is where the individual-facing description axis touches the ensemble-facing channel axis.

4.5 Solomonoff induction

The universal distribution also founds a theory of prediction. Solomonoff's universal induction [10] predicts the continuation of a sequence by conditioning the universal distribution on the observed prefix, weighting every computable hypothesis by $2^{-(\text{its length})}$ — a formal rendering of Occam's razor in which shorter descriptions are a priori more probable. On the substrate reading, Solomonoff induction is the description axis turned toward prediction: it uses the shortest-description measure not to score a fixed individual but to extrapolate an unfinished one, and its celebrated optimality is the statement that no computable predictor does essentially better than weighting by description length. We locate it as the predictive employment of the description axis and return, in Section 6, to the uncomputability it shares with the measure it is built on.

4.6 Counterarguments and Responses

Objection 1 (the logarithmic terms make the algorithmic apparatus messy compared to Shannon's clean identities). *Response.* They make it inexact, not messy, and the inexactness is informative rather than regrettable. The exact additivity of the channel axis and the up-to-log additivity of the description axis are different structural facts about two different reductions, and Section 5 shows that the difference is exactly what determines the form of the morphism between the two axes. A theory that hid the logarithmic terms would hide the feature that distinguishes the description axis from the channel axis.

Objection 2 (the coding theorem already identifies K with a probability, so the description axis just is a channel axis in disguise). *Response.* It identifies the description *measure* of an individual with the negative log of a particular distribution, but that distribution — the universal distribution — is itself defined from description lengths, not from a source's repetition statistics. The coding theorem is a relation *between* the axes, the per-object averaging morphism, not a collapse of one into the other. The description axis remains the axis of the individual; the coding theorem says how, averaged, it meets the axis of the ensemble.

4.7 Section Result and Implications

We have exhibited the algorithmic apparatus — prefix complexity, conditional and joint complexity with the chain rule, algorithmic mutual information with its symmetry (Proposition 4.2), the universal distribution and Levin's coding theorem (Proposition 4.3), and Solomonoff induction — as constructs on the retained individual. The implication is that the description axis, like the channel axis, carries not a bare functional but a complete theory, here a theory of the individual object and its prediction. We now describe the shape of that axis.

5 The Shape of the Axis

5.1 Native indeterminacy and composition law

Two structural features parallel, and contrast with, those of the channel axis. The *native indeterminacy* of the description axis is the additive machine constant of Proposition 4.1:

no fact about the individual fixes the choice of universal machine, and different machines shift every complexity by a bounded amount. This is the description-axis analogue of the channel axis's free unit, with one difference worth noting: the channel unit is a multiplicative scale (a choice of logarithm base), while the machine constant is an additive offset. The *composition law* is where the axes diverge most sharply.

Proposition 5.1 (Subadditivity, not exact additivity)**.** *For all x, y,*

$$K(x, y) \leq K(x) + K(y) + O(\log), \tag{4}$$

and the inequality is in general strict: there is no exact additivity $K(x, y) = K(x)+K(y)$ even for "independent" individuals, because the description must encode, at logarithmic cost, where x ends and y begins. The description axis is subadditive up to a logarithmic term, not exactly additive.

Proof. Concatenating shortest prefix programs for x and y yields a program for the pair; the prefix property allows the decoder to separate them at an additional cost logarithmic in the lengths, giving (4) [8]. Strictness follows because shared structure between x and y reduces $K(x, y)$ below the sum, and even structureless pairs incur the logarithmic separation cost; no choice of machine removes it for all pairs. □

This is the structural fact behind a feature of Article 3's typology that might otherwise look arbitrary. The channel axis, exactly additive, enters a *rescaling* rigidity with the macrostate axis: thermodynamic entropy is k_B times a channel entropy, a proportionality, and exact additivity is what a proportionality requires. The description axis, additive only up to log, cannot enter such a proportionality with the channel axis; it enters instead an *averaging* relation, $\mathbb{E}[K] = H + O(1)$, in which the channel entropy is the expectation of the description complexity over a source. The kernel morphism of Article 3 — averaging the description projection over an ensemble to recover the channel projection — is the only relation the inexact additivity permits, and Levin's coding theorem (Proposition 4.3) is its per-object statement. The shape of the axis thus predicts the shape of its relation to the channel axis: subadditive measures average, they do not rescale.

5.2 Randomness as incompressibility

The description axis gives randomness its sharpest definition: an individual is random when it has no shorter description than itself.

Proposition 5.2 (Incompressibility of the typical individual)**.** *Among the 2^n binary strings of length n, fewer than 2^{n-k} have $K(x) < n - k$; hence the fraction of n-bit strings compressible by more than k bits is less than 2^{-k}. Almost every individual is incompressible: for most x of length n, $K(x) \geq n - O(1)$.*

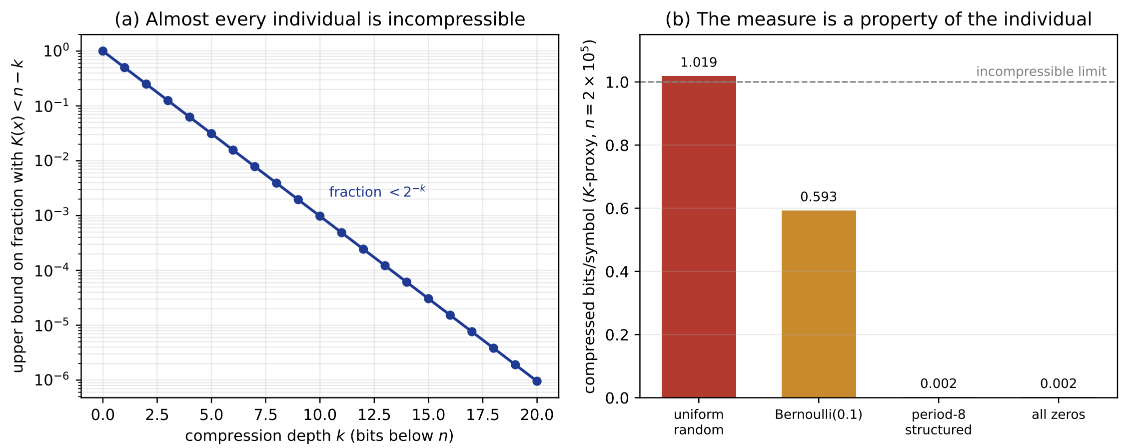

Figure 1: Randomness as incompressibility on the description axis. **(a)** The counting bound of Proposition 5.2: the fraction of n-bit individuals compressible by more than k bits is below 2^{-k} and collapses exponentially, so almost every individual is incompressible — the description axis cannot shorten the typical object. **(b)** The measure is a property of the individual. Using a real lossless compressor (bz2) as a computable *upper* proxy for K, the compressed bits per symbol of four n-bit individuals: the i.i.d. uniform-random string is incompressible (proxy near the 1-bit ceiling), the biased and periodic strings are progressively compressible, and the all-zeros string is maximally compressible. The description-axis measure is maximal for the random individual and small for the structured one, the signature of an axis that retains the individual rather than the ensemble.

Proof. There are at most $2^{n-k} - 1$ programs of length less than $n - k$, so at most that many strings have a description shorter than $n - k$; dividing by 2^n gives a fraction below 2^{-k} [8]. The bound is independent of n. □

On the substrate reading, the random individual is the one the description reduction retains in full but can compress not at all: it is all individual and no regularity, so its shortest description is itself. This is the opposite pole from the channel axis's uniform source, which maximizes entropy by maximizing the ensemble's spread; here maximality is a property of a single object, its incompressibility. The two notions of "maximal information" — maximal entropy on the channel axis, maximal complexity on the description axis — are maxima of different reductions, and they coincide only on average, through the coding theorem. Figure 1 shows both the exact counting bound and an empirical confirmation that real compressors leave the random individual essentially uncompressed while collapsing the structured ones.

5.3 Counterarguments and Responses

Objection 1 (using bz2 as a proxy for K is illegitimate, since K is uncomputable). *Response.* The figure uses bz2 only as a computable *upper bound* on K, and labels it as such. Any lossless compressor furnishes an upper proxy — it exhibits one description, not the shortest — so its output bounds K from above and can demonstrate compressibility, though never incompressibility, of a given object. The incompressibility claim of panel (a) is the exact counting bound of Proposition 5.2, not an empirical compressor result; panel (b) illustrates only the uncontroversial direction, that structured individuals admit short descriptions.

Objection 2 (if the description axis is subadditive and the channel axis additive, one of them is the wrong measure). *Response.* Neither is wrong; they measure different reductions, and the composition law of each is correct for what it measures. Exact additivity is correct for an ensemble, where independent sources contribute independently; subadditivity up to log is correct for individuals, where concatenation incurs a separation cost and shared structure discounts the sum. The difference is a fact about the reductions, and Section 5.1 shows it to be exactly what fixes the morphism between them.

5.4 Section Result and Implications

We have exhibited the shape of the description axis: its native indeterminacy is the additive machine constant, its composition law is subadditivity up to a logarithmic term (Proposition 5.1), and that inexact additivity is the structural reason the channel–description relation is a kernel (averaging) morphism rather than a rescaling. Randomness is located as incompressibility (Proposition 5.2, Figure 1). The implication is that the form of the axis predicts the form of its relations. We turn to its characteristic limitation.

6 Uncomputability

6.1 The price of retaining the whole individual

The channel axis's characteristic feature was observer-relativity, traced to the dependence of the channel reduction on a prior choice. The description axis's characteristic feature is uncomputability, and we trace it, likewise, to the defining choice of the axis.

Proposition 6.1 (K is upper semi-computable but not computable)**.** *The prefix complexity K is upper semi-computable: there is an algorithm that, given x, produces a non-increasing sequence converging to $K(x)$ from above (by running programs and recording the shortest that has halted so far). But K is not computable: no algorithm, given x, outputs $K(x)$ exactly and halts for all x [8].*

Proof. Upper semi-computability is the dovetailed search over programs, which can only lower the current best bound. Non-computability follows from the halting problem: a total algorithm for K would let one compute, for each n, the lexicographically first string of complexity at least n, an object of complexity $\geq n$ produced by a program of length $O(\log n)$ — a contradiction for large n (a form of the Berry paradox made rigorous) [2, 8]. □

The uncomputability is the price of the axis's defining choice. The description reduction retains the whole individual (Definition 3.1), and the shortest specification of a whole individual is not effectively findable: to know it one would have to rule out every shorter program, which requires deciding which programs halt. An axis that retains the entire individual must, by the halting problem, leave the minimal specification of that individual uncomputable. This is not a flaw in Kolmogorov's definition but a structural consequence of measuring the individual rather than the ensemble: the channel axis, measuring only repetition statistics, is computable from those statistics; the description axis, measuring the individual, inherits the individual's full intractability.

6.2 Uncomputability bears on evaluation, not derivation

It is essential, and was already noted in Article 3, that the uncomputability of K bears on the *evaluation* of the measure, not on its *derivation* or its standing as the description-axis projection. That K is the shortest-description measure, fixed up to a constant by invariance, is a theorem; that its value for a given x cannot be computed is a separate fact about evaluability. The substrate reading respects the distinction: the description axis is well-defined as a projection — the relations it enters with the other axes (the averaging rigidity, the refinement decomposition) are exact theorems — even though the projection's value on a particular individual is not algorithmically accessible. A measure can be derived, characterized, and related to others while remaining uncomputable; algorithmic information theory is the standing proof that it can.

Remark 6.2 (Uncomputability as the inverse of a lossy projection)**.** There is a reading of the uncomputability that connects it to the asymmetric bridge of the volume. Recovering the shortest description of an individual is the inverse problem of a lossy projection: many programs produce the same output, the production map is many-to-one and not effectively invertible, and finding the minimal preimage is exactly the search the halting problem obstructs. On this reading the uncomputability of K is the description-axis instance of the theme that projection loses and the model must recover by hypothesis (Part III): a subsystem that would compute K would have to invert a lossy map, and cannot do so effectively. The channel axis met this asymmetry as the data-processing inequality [6]; the description axis meets it as uncomputability. Both are faces of the same lossiness, on different axes.

6.3 Counterarguments and Responses

Objection 1 (an uncomputable measure is useless, so the description axis is idle). *Response.* It is uncomputable but not useless, and the two are routinely confused. Upper proxies (any compressor) furnish computable bounds; the measure's exact relations to other axes are theorems that hold without evaluating it; and Solomonoff induction's optimality is a statement about the measure, not a procedure requiring its computation. Much of mathematics concerns objects that cannot be exhibited algorithmically; the description axis is a well-defined projection whose value is not computable, which constrains its use without nullifying its theory.

Objection 2 (calling uncomputability "the characteristic limitation" overstates a technicality). *Response.* It is not a technicality but the structural counterpart, on the description axis, of the channel axis's observer-relativity: each axis has a defining choice (the source model; the universal machine) and a characteristic limitation that flows from what it retains (relativity to the chosen reduction; uncomputability of the minimal description). Presenting uncomputability as the description axis's signature feature is the parallel the classification requires, and it earns its place by the connection of Remark 6.2 to the volume's central asymmetry.

6.4 Section Result and Implications

We have located uncomputability as the price of the description axis's defining choice — K is upper semi-computable but not computable (Proposition 6.1), because the minimal specification of a retained individual is not effectively findable — and we have been careful that this bears on evaluation, not on the derivation of the measure or its relations (Section 6.2), with the uncomputability read as the inverse of a lossy projection (Remark 6.2). The implication is that the description axis, like the channel axis, carries a characteristic limitation that the substrate reading locates rather than laments. We close by marking what the axis cannot see.

7 What the Description Axis Cannot See

The description axis is a complete theory of the individual object, and its completeness is real: within the axis, complexity, conditional and joint complexity, algorithmic mutual information, the universal distribution, and Solomonoff induction answer the questions the axis can pose. But a theory that takes the description measure as primitive is blind, from within, to the relations between the description projection and the others.

It cannot see why the expected shortest description length of a source equals its entropy — the agreement with the channel axis — because the source does not appear in it; the individual $K(x)$ knows nothing of the ensemble whose expectation recovers H. That agreement is the averaging rigidity of Article 3, exact at the rate level by Brudno's theorem [1, 12], and visible only from the substrate, where the description projection and the channel projection are reductions of one source. And it cannot see why an individual's complexity splits into a structural part and a random residue — the relation to the regularity axis — because the regularity axis does not appear in it either; the two-part code $K(x) \approx K(\text{model}) + (\text{bits given the model})$, and the minimal sufficient statistic that realizes it [11], are the refinement decomposition of Article 3, developed for the regularity axis in Article 7 of this volume. The description axis records the total; it does not, from within, split it.

This is the same sense in which the substrate frame held Shannon's theory: it assigns the algorithmic theory its facet — the description axis, with its individual-retaining reduction, its prefix functional, its universal distribution — and affirms its completeness there, then supplies what the facet cannot, the relations to the other facets as the rigidities of the morphisms between reductions. We locate the theory's authors with the generosity the volume owes its predecessors. Solomonoff [10] founded the algorithmic theory of induction and the universal prior; Kolmogorov [3] gave the complexity of the individual its definition and the invariance that makes it well-posed; Chaitin [2] supplied the prefix formulation that aligns it with information theory; Levin [7] proved the coding theorem that links it to probability. Each saw the description facet truly, and built the theory that the substrate frame now places beside the channel theory as the second complete account of a projection.

7.1 Counterarguments and Responses

Objection 1 (if the description axis is complete within itself, the substrate adds nothing a working theorist needs). *Response.* It adds nothing *within* the description axis, which is what completeness means; it adds what is needed at the borders — relating a complexity to an entropy through the coding theorem and the averaging rigidity, or splitting a complexity into structure and residue through the refinement decomposition. Those cross-axis relations are where a single-axis theory reports coincidences and where the substrate frame supplies derivations (Article 3). The addition is at the borders, not within.

Objection 2 (the two-part code belongs to algorithmic information theory already, so the regularity relation is not the substrate's to claim). *Response.* The two-part code and the minimal sufficient statistic are indeed algorithmic constructions [11], and we claim no priority over them. What the substrate frame contributes is the reading of the decomposition as a *refinement morphism* between the description axis and the regularity axis — one of the classified inter-axis maps of Article 3 — which locates the construction within a typology and connects it to the rescaling and averaging morphisms of the other axes. The construction is theirs; its place in the morphism typology is the frame's.

7.2 Section Result and Implications

We have marked what the description axis cannot see — the averaging relation to the channel axis and the refinement decomposition toward the regularity axis, both supplied by Article 3 — and located Kolmogorov, Solomonoff, Chaitin, and Levin as the authors of the complete theory of the description facet. The implication, carried forward in Part II, is that the regularity axis, whose decomposition the description axis records but cannot split, receives its own located treatment next, in the article on effective complexity.

8 Epistemic Status: Proved, Imported, Affirmed

In the manner of the series, we close with an explicit ledger. *Proved (here)* marks results established in this article; *Imported* marks results taken from prior articles, chiefly the derivation and the cross-axis rigidities of Article 3; *Affirmed (standard)* marks established algorithmic information theory the article endorses and locates without re-deriving.

Claim	**Status**	**Locus**
K = description-axis projection, unique up to constant	Imported	Kriger [5] (Art. 3)
Averaging rigidity $\mathbb{E}[K] = H+O(1)$; refinement decomposition	Imported	Kriger [5] (Art. 3)
Description reduction retains the individual, discards the ensemble	Proved (here)	Def. 3.1
Orthogonality of description and channel reductions	Proved (here)	Rem. 3.2
Symmetry of algorithmic information	Proved (here)	Prop. 4.2
Subadditivity up to log, not exact additivity	Proved (here)	Prop. 5.1
Inexact additivity $\Rightarrow$ kernel morphism, not rescaling	Proved (here)	§5.1
Incompressibility of the typical individual	Proved (here)	Prop. 5.2
K upper semi-computable but not computable	Proved (here)	Prop. 6.1
Uncomputability bears on evaluation, not derivation	Proved (here)	§6.2
Prefix complexity; invariance; chain rule	Affirmed (standard)	Chaitin [2], Li and Vitányi [8]
Levin's coding theorem; universal distribution	Affirmed (standard)	Levin [7]
Solomonoff induction and its optimality	Affirmed (standard)	Solomonoff [10]
Algorithmic randomness (Martin-Löf)	Affirmed (standard)	Martin-Löf [9]
Two-part code / minimal sufficient statistic	Affirmed (standard)	Vereshchagin and Vitányi [11]

The ledger marks the division of labor. The article *imports* the derivation and the cross-axis rigidities from Article 3; it *proves* the structural facts about the description reduction, the apparatus built on it, the shape of the axis, and its uncomputability; and it *affirms*, without re-deriving, the established algorithmic information theory it locates.

9 Discussion

The result of this article is a placement, the second of Part II. Kolmogorov complexity is located on the description axis as the complete theory of the individual object: the description reduction retains the whole individual and discards the ensemble, the prefix functional measures the shortest specification, and the apparatus of conditional and algorithmic mutual information, the universal distribution, and Solomonoff induction are constructs on the retained individual. The placement makes two contributions beyond the bookkeeping, paralleling the channel article's two. It shows that the inexact additivity of the description axis — subadditivity up to a logarithmic term, against the exact additivity of the channel axis — is not a blemish but the structural reason the two axes relate by averaging rather than rescaling, with Levin's coding theorem the per-object statement of the average. And it locates the uncomputability of K as the price of retaining the whole individual, the description axis's characteristic limitation, connected through the inverse-of-a-lossy-projection reading to the asymmetry the volume develops elsewhere.

The placement continues the discipline of the series. It does not demote the algorithmic theory; it affirms its completeness within its axis and reserves only the cross-axis relations for the substrate. It does not promote the description measure to the ground, as the reification pictures of Article 2 would; the shortest description of an individual is a description-axis quantity, and the individual a projection, never the substrate. And it marks honestly the limit of the axis: a theory that takes the description measure as primitive cannot see why its measure agrees, on average, with the channel measure, nor how it splits relative to the regularity measure, because the common domain that induces both relations is exactly what a single-axis theory omits.

9.1 Limitations

Three limitations are material. *First*, the article imports the derivation of K as the description-axis projection and the cross-axis rigidities from Article 3 and does not re-establish them; a reader who doubts those results will find the placement only as secure as its source. *Second*, the empirical panel of Figure 1 uses a real compressor as a computable upper proxy for K and can therefore demonstrate compressibility but never incompressibility; the incompressibility claim rests on the exact counting bound of Proposition 5.2, not on the compressor. *Third*, the refinement decomposition toward the regularity axis is set up here but developed only in Article 7; the present article records that the description axis carries the total complexity without, from within, splitting it into structure and residue.

9.2 Future Directions

The immediate sequel is Article 7, which does for the regularity axis of effective complexity what this article and its predecessor did for the description and channel axes, and which develops the two-part code the description axis here only records. Beyond the volume, the contrast between the exactly additive channel axis and the up-to-log subadditive description axis invites a precise question continuous with Article 3: whether the logarithmic slack of the description axis can be characterized as a controlled deformation of an otherwise additive law — a deformation morphism, in the sense raised for the Rényi and Tsallis families on the channel axis — and whether the typology of inter-axis maps should include such within-axis deformations alongside the rescaling, kernel, and refinement morphisms between axes. That question connects the present located treatment to the open completeness conjecture of Article 3 and is left to the work that can address it.

Acknowledgments

The author thanks colleagues of the Information Physics Institute and the Institute of Integrative and Interdisciplinary Research for sustained discussion of the foundations of the series.

References

[1] A. A. Brudno. Entropy and the complexity of the trajectories of a dynamical system. *Transactions of the Moscow Mathematical Society*, 44:127–151, 1983.

[2] G. J. Chaitin. A theory of program size formally identical to information theory. *Journal of the ACM*, 22(3):329–340, 1975.

[3] A. N. Kolmogorov. Three approaches to the quantitative definition of information. *Problems of Information Transmission*, 1(1):1–7, 1965.

[4] B. Kriger. The information substrate: differentiation as the ground of information. The Information Substrate Theory (IST), Volume VII, Article 1. Zenodo, 2026.

[5] B. Kriger. The projection theorem: how known information measures descend from the substrate. The Information Substrate Theory (IST), Volume VII, Article 3. Zenodo, 2026.

[6] B. Kriger. The axis of the channel: Shannon information as a substrate projection. The Information Substrate Theory (IST), Volume VII, Article 5. Zenodo, 2026.

[7] L. A. Levin. Laws of information conservation (nongrowth) and aspects of the foundation of probability theory. *Problems of Information Transmission*, 10(3):206–210, 1974.

[8] M. Li and P. Vitányi. *An Introduction to Kolmogorov Complexity and Its Applications*. Springer, Cham, 4th edition, 2019.

[9] P. Martin-Löf. The definition of random sequences. *Information and Control*, 9(6):602–619, 1966.

[10] R. J. Solomonoff. A formal theory of inductive inference, parts I and II. *Information and Control*, 7(1):1–22 and 7(2):224–254, 1964.

[11] N. K. Vereshchagin and P. M. B. Vitányi. Kolmogorov's structure functions and model selection. *IEEE Transactions on Information Theory*, 50(12):3265–3290, 2004.

[12] A. K. Zvonkin and L. A. Levin. The complexity of finite objects and the development of the concepts of information and randomness by means of the theory of algorithms. *Russian Mathematical Surveys*, 25(6):83–124, 1970.

A Code for Figure 1

The following Python script generates Figure 1. Panel (a) plots the exact counting bound 2^{-k} of Proposition 5.2; panel (b) uses bz2 as a computable upper proxy for K on four n-bit individuals.

```
import bz2
import numpy as np
rng = np.random.default_rng(20260531)

def packed(bits):
    pad = (-len(bits)) % 8
    if pad: bits = np.concatenate([bits, np.zeros(pad, dtype=np.uint8)])
    return np.packbits(bits).tobytes()

def proxy_bits_per_symbol(bits):           # computable UPPER bound on K/n
    return 8.0 * len(bz2.compress(packed(bits), 9)) / len(bits)

ks = np.arange(0, 21)
bound = 2.0 ** (-ks)                        # fraction with K(x) < n-k  is  < 2^{-k}

n = 200_000
uniform  = (rng.random(n) < 0.5).astype(np.uint8)      # ~ 1.019  (incompressible)
biased   = (rng.random(n) < 0.1).astype(np.uint8)      # ~ 0.593
periodic = np.tile([1,1,0,1,0,0,0,1], n//8 + 1)[:n]    # ~ 0.002  (period-8)
zeros    = np.zeros(n, dtype=np.uint8)                 # ~ 0.002  (all zeros)
```

The Axis of Regularity:
Effective Complexity and the Structured Projection

The complete theory of structure, located: the two-part code, the non-monotonic signature, and the relativity of the cut

Boris Kriger[1,2]

[1] Information Physics Institute
boris.kriger@informationphysicsinstitute.net

[2] Institute of Integrative and Interdisciplinary Research, Toronto
boriskriger@interdisciplinary-institute.org

ORCID: https://orcid.org/0009-0001-0034-2903

The Information Substrate Theory (IST) — Volume VII, Article 7
Part II: The Projections — classification of information theories

Abstract

Shannon's entropy and Kolmogorov's complexity both rise with disorder: each is small for a perfectly ordered object and large for a perfectly random one. Effective complexity does neither. It is small at both extremes — the ordered object has little structure to capture, the random object has none — and large only in between, for the structured object whose regularities are themselves substantial. This article, the third of the classificatory part of the volume, locates effective complexity on the regularity axis of the substrate as the complete theory of structure. The regularity reduction is a refinement of the description reduction: it splits the shortest description of an individual into a model of its regularities and a random residue, and retains only the model. The Projection Theorem of Article 3 established this split, $K = K_{\text{eff}} + R$, as the refinement morphism between the description and regularity axes; here we develop the axis in full. We exhibit its apparatus — the two-part code, the Gell-Mann–Lloyd total information, the Kolmogorov structure function, and the algorithmic minimal sufficient statistic — as constructs on the retained model, and we locate Bennett's logical depth as a neighboring regularity-type measure that shares the non-monotonic signature while measuring a different resource. The signature itself, the non-monotonic "hump," is given its precise statement and a concrete measured demonstration: structured and random objects of nearly equal total complexity can have sharply different effective complexity, so effective complexity is not a function of the description-axis or channel-axis measure alone. Finally we locate the characteristic feature of the axis — the relativity of the structure/residue cut — as the origin of effective complexity's reputation for contestedness, the regularity axis's own native indeterminacy, the structural counterpart of the channel axis's observer-relativity and the description axis's uncomputability.

Keywords: effective complexity; regularity axis; two-part code; Kolmogorov structure function; minimal sufficient statistic; logical depth; total information; non-monotonic complexity; foundations of physics.

Contents

1 Introduction

The two preceding articles located the channel and description axes: Shannon's entropy as the registered differentiation per draw of a source-as-ensemble [6], and Kolmogorov complexity as the shortest specification of an individual-as-object [7]. Both measures share a feature that, on reflection, is a limitation as much as a virtue: they are monotone in disorder. Entropy is least for a settled source and greatest for a uniform one; complexity is least for a compressible string and greatest for an incompressible one. Neither distinguishes the rich from the random. A library and a shelf of noise can have the same entropy; a working program and a random bit-string of the same length can have the same Kolmogorov complexity. What both measures miss is exactly what we usually mean by complexity in the ordinary sense — not sheer disorder, but *structure*: the property an object has when it is neither trivially simple nor merely random, but organized.

Effective complexity is the measure built to capture this [3]. Its defining feature is that it is *non-monotonic* in disorder: small for the ordered, small for the random, large only for the structured object in between. This article locates effective complexity on the regularity axis of the substrate, as the complete theory of structure, and shows that its non-monotonicity — often treated as a puzzle or a defect — is the exact signature of the reduction it rests on.

The placement is established in outline. The Projection Theorem of Article 3 [5] identified the regularity axis as a *refinement* of the description axis. The shortest description of an individual splits, by the two-part code, into a model of its regularities and a random residue: $K(x) = K_{\text{eff}}(x) + R(x)$, with K_{eff} the complexity of the model and R the residual randomness. The description axis retains the total K; the regularity axis retains only the model K_{eff}, discarding the residue R. This refinement morphism — an additive split, with no proportionality — is what relates the two axes, and it is why effective complexity is a genuinely distinct projection rather than a rescaling of complexity. The present article develops the axis in the detail a dedicated treatment permits, exactly as the channel and description articles developed theirs.

Two contributions go beyond the bare placement, paralleling those of the companion articles. The first concerns the signature. We give the non-monotonicity its precise statement — that K_{eff} vanishes, up to constants, at both the ordered and the random extreme — and a concrete measured demonstration that structured and random objects of nearly equal total complexity can have sharply different effective complexity, so that effective complexity is not a function of the description-axis measure alone. The second concerns the characteristic feature of the axis. Effective complexity is famously contested: critics observe that what counts as "regularity" rather than "residue" depends on a prior judgment, and conclude that the measure is subjective. We locate this dependence precisely. It is the relativity of the *cut* between structure and residue — the level at which the structure function is read — and it is the regularity axis's own native indeterminacy, the structural counterpart of the channel axis's observer-relativity and the description axis's uncomputability. Given the cut, effective complexity is determined; the cut is the choice the reduction requires, and locating it dissolves the charge of subjectivity into a precise account of what must be fixed.

We close, as before, by marking what the regularity axis cannot see and locating its authors generously. Gell-Mann and Lloyd built the theory of structure; Bennett built a neighboring one; the substrate frame assigns each its facet and explains why the facet relates to the others as it does.

2 Logical Structure of the Argument

The argument parallels those of the channel and description articles, for the third axis.

Step 1 (Section 3): The regularity reduction.
The regularity reduction refines the description reduction, splitting the shortest description into a model and a residue and retaining only the model (Definition 3.1); effective complexity is the complexity of the model.

Step 2 (Section 4): The apparatus.
We exhibit the two-part code, the Gell-Mann–Lloyd total information, the Kolmogorov structure function, and the algorithmic minimal sufficient statistic (Definition 4.2) as constructs on the retained model, recalling the refinement morphism of Article 3 (Proposition 4.1).

Step 3 (Section 5): The shape of the axis.
The native indeterminacy is the machine constant together with the structure/residue cut; the signature is the non-monotonic hump (Proposition 5.1, Figure 1); and effective complexity is not a function of the description- or channel-axis measure alone (Proposition 5.2). Bennett's logical depth is located as a neighboring regularity-type measure.

Step 4 (Section 6): The characteristic feature.
Effective complexity is determined only relative to the structure/residue cut (Proposition 6.1); this is the regularity axis's native indeterminacy, and locating it dissolves the charge of subjectivity.

Step 5 (Section 7): What the regularity axis cannot see.
The axis is the theory of structure and is blind, from within, to the common domain that fixes the cut; we locate Gell-Mann, Lloyd, and Bennett generously.

Step 6 (Section 8): Epistemic status.
An explicit ledger separates what is proved here, what is imported, and what is affirmed as standard.

The dependencies are linear: Steps 2–4 operate on the reduction of Step 1; Step 5 sets the whole within the cross-axis frame of Article 3. The thesis advanced is that effective complexity is the complete and correct theory of the structured facet, and that its non-monotonic signature and its contested cut are exact features of the regularity reduction rather than defects.

3 The Regularity Reduction

3.1 Structure retained, residue discarded

The regularity axis stands to the description axis not as an orthogonal coordinate, as the channel axis did, but as a refinement: it takes the shortest description and keeps only its structured part.

Definition 3.1 (Regularity reduction)**.** The *regularity reduction* ρ_{re} acts on the description projection by splitting the shortest description of an individual x into two parts — a *model* M capturing the regularities of x, and a *residue* specifying x among the objects the model does not distinguish — and retaining only the model. The regularity projection is the model M, and the measure on it, *effective complexity* $K_{\mathrm{eff}}(x)$, is the complexity $K(M)$ of the model in the optimal such split.

What the regularity reduction discards is the residue: the part of the description that is specific to the individual and carries no regularity, the incompressible remainder once the structure has been factored out. What it retains is the model: the description of the regularities alone. This is exactly the complement of what the description axis discards. The description axis retained the whole individual — structure and residue together — and discarded the ensemble; the regularity axis goes one step further, retaining only the structure within the individual and discarding the residue as well. The relation between them is therefore the refinement morphism of Article 3, made precise in Section 4.

Remark 3.2 (Why "refinement" and not "orthogonality")**.** The channel and description axes were orthogonal: each retained what the other discarded, and neither was a further reduction of the other. The regularity axis is not orthogonal to the description axis but *downstream* of it: it operates on the description, splitting it. This is why Article 3 classified the description–regularity relation as a refinement morphism rather than a kernel or a rescaling, and why effective complexity is built from Kolmogorov complexity rather than defined independently. The regularity axis is the description axis seen through the structure/residue split.

3.2 The structured object as model-bearer

To take the regularity reduction is to treat the object as a *model-bearer*: something whose regularities can be separated from its accidents and described on their own. This presupposes a prior judgment of which features are regularities and which are residue — the cut of Section 6 — just as the channel reduction presupposed a source model and the description reduction a universal machine. The cut is the structure the regularity reduction requires, and effective complexity is a property of the object relative to it, never of the substrate (Article 2): the substrate has no model and no residue; the split is a feature of the reduction.

3.3 Counterarguments and Responses

Objection 1 (effective complexity is just Kolmogorov complexity with extra steps). *Response.* It is built from Kolmogorov complexity but is not a function of it: two objects of equal K can have different K_{eff} (Proposition 5.2), because the split of K into model and residue differs between them. The regularity reduction extracts information the description projection contains but does not isolate — the share of the description that is structural — and that extraction is a genuine further reduction, not a relabeling.

Objection 2 (splitting a description into "model" and "residue" is arbitrary). *Response.* The split is not arbitrary once the cut is fixed; it is the optimal two-part code, made precise by the Kolmogorov structure function and the minimal sufficient statistic (Section 4). What is genuinely chosen is the cut — the level at which structure gives way to residue — and that choice is the axis's native indeterminacy, located in Section 6, not an arbitrariness in the split given the cut.

3.4 Section Result and Implications

We have fixed the regularity reduction as the refinement of the description reduction that retains the model and discards the residue (Definition 3.1, Remark 3.2). The implication is that every quantity on the regularity projection is a property of the structure of an object, separated from its randomness — a profile the two-part code and the structure function make precise, as we now show.

4 The Apparatus

4.1 The refinement morphism, recalled

The relation that defines the regularity axis is the two-part split of the description, recalled from Article 3.

Proposition 4.1 (Effective complexity as regularity-axis projection; 5)**.** *The shortest description of an individual x admits a two-part form*

$$K(x) = K(M) + K(x \mid M) + O(1), \tag{1}$$

where M is a model (a finite set, or program for a set, containing x) and $K(x \mid M)$ is the cost of specifying x within M. For the optimal model — the algorithmic minimal *sufficient statistic — the second term is the residual randomness $R(x)$ and the first is the effective complexity $K_{\text{eff}}(x) = K(M)$, so $K(x) = K_{\text{eff}}(x) + R(x) + O(1)$. Effective complexity is the regularity-axis projection $K_{\text{eff}} \circ \rho_{\text{re}}$, and the split is the refinement morphism between the description and regularity axes.*

The split is additive and carries no proportionality: K_{eff} is not a fixed fraction of K, nor a rescaling of it, but the model-share of the description, which varies independently of the total (Section 5). This is the structural reason the description–regularity relation is a refinement morphism and not a rescaling like the channel–macrostate relation: the two-part code adds, it does not scale.

4.2 Total information, the structure function, and the minimal sufficient statistic

The apparatus of the regularity axis makes the optimal split precise.

Definition 4.2 (The regularity-axis constructs)**.** For an individual x of length n:

- The *Gell-Mann–Lloyd total information* is $\Sigma(x) = K_{\text{eff}}(x) + H(M)$, the effective complexity plus the entropy of the optimal model's set — the structural bits plus the residual bits [3, 4].
- The *Kolmogorov structure function* is $h_x(\alpha) = \min\{\log|M| : x \in M,\ K(M) \leq \alpha\}$, the least log-cardinality of a model of complexity at most α containing x [2, 8].
- The *algorithmic minimal sufficient statistic* is the smallest model M at which $h_x(\alpha)+\alpha$ reaches $K(x)$ — the point where adding model complexity no longer reduces the residue. Its complexity is $K_{\text{eff}}(x)$.

The structure function is the rigorous backbone of effective complexity. Reading it from left to right, one spends increasing complexity α on the model and watches the residual $\log|M|$ fall; the minimal sufficient statistic is the "elbow" where the residual stops falling, beyond which extra model complexity merely memorizes the random residue and buys nothing. Effective complexity is the location of that elbow: the model complexity at which all the structure, and none of the randomness, has been captured. The total information Σ adds the residual back, recovering a quantity close to the total description length; effective complexity alone is the structural part.

4.3 Logical depth: a neighboring regularity-type measure

Effective complexity is not the only measure built to capture structure rather than disorder, and locating its neighbor sharpens the axis. Bennett's *logical depth* [1] measures the structure of an object by the *time* its shortest program takes to produce it, rather than the *size* of its model: a deep object is one whose production requires long computation from a short description, so that its organization is the buried result of much work. Logical depth shares effective complexity's non-monotonic signature — both the ordered and the random object are shallow, the first because its program is short and fast, the second because its shortest program is essentially itself and runs at once — but it measures a different resource, computation time rather than model size. On the substrate reading, logical depth is a neighboring regularity-type axis, related to effective complexity as the Rényi and Tsallis families were related to Shannon on the channel axis: a different functional on a kindred reduction, sharing the structured-not-random signature, differing in the resource it counts. We locate it as such and do not adjudicate which captures "structure" best; that depends on whether size or time is the structure of interest.

4.4 Counterarguments and Responses

Objection 1 (the structure function and minimal sufficient statistic are uncomputable, so the apparatus is empty). *Response.* They inherit the uncomputability of K [7], and the response is the same: uncomputability bears on evaluation, not on definition or on the relations the axis enters. The structure function is a well-defined object whose shape is the subject of theorems [8], computable upper proxies exist (any explicit two-part code furnishes one, as in Figure 1), and the refinement morphism relating effective complexity to the description axis is an exact relation. An uncomputable measure with exact relations and computable proxies is not empty.

Objection 2 (logical depth and effective complexity cannot both be "the" measure of structure). *Response.* Neither is "the" measure; they are two functionals on the regularity facet, sharing its signature and differing in resource. The substrate frame does not crown one; it locates both on the regularity axis and notes that the choice between size and time is a choice of what structure one means, exactly as the choice among the Rényi orders was a choice of composition law on the channel axis. Plurality within an axis is located, not adjudicated.

4.5 Section Result and Implications

We have exhibited the apparatus of the regularity axis — the two-part code and refinement morphism (Proposition 4.1), the total information, structure function, and minimal sufficient statistic (Definition 4.2), and the neighboring measure of logical depth — as constructs on the retained model. The implication is that the regularity axis, like its predecessors, carries a complete apparatus, here a theory of the structured part of an object. We turn to the shape of the axis, where its signature lives.

5 The Shape of the Axis

5.1 Native indeterminacy

The regularity axis carries two sources of native indeterminacy, where the channel and description axes carried one each. It inherits the additive *machine constant* from the Kolmogorov complexity on which it is built (Proposition 4.1). And it carries, additionally, the *structure/residue cut*: the level at which the structure function is read, the boundary between what counts as model and what counts as residue. The cut is the regularity axis's own indeterminacy, with no analogue on the channel or description axes, and it is the subject of Section 6. Given both — a machine and a cut — effective complexity is determined; the two are the structure the reduction requires.

5.2 The non-monotonic signature

The defining feature of the regularity axis is the shape of its measure across the order–randomness range.

Proposition 5.1 (Non-monotonicity of effective complexity)**.** *Effective complexity vanishes, up to additive constants, at both extremes of disorder:*

(i) (Ordered.) *If x is highly compressible, $K(x) = O(\log n)$, then $K_{\text{eff}}(x) = O(\log n)$: there is little description to split, so the model is small.*

(ii) (Random.) *If x is incompressible, $K(x) = n + O(\log n)$, then $K_{\text{eff}}(x) = O(\log n)$: the minimal sufficient statistic is trivial, the residue is essentially all of x, and the model is small.*

Effective complexity is large only for intermediate, structured x, whose regularities are themselves substantial. It is therefore non-monotonic in disorder, unlike Shannon entropy and Kolmogorov complexity, which are monotone.

Proof. (i) If $K(x) = O(\log n)$ then by $K(x) = K_{\text{eff}}(x)+R(x)+O(1)$ and $K_{\text{eff}}, R \geq 0$ we have $K_{\text{eff}}(x) \leq K(x) = O(\log n)$. *(ii)* For an incompressible x, the structure function h_x lies near the diagonal: no model of complexity below $K(x)$ reduces the residue below $\log |M|$, so the minimal sufficient statistic is the trivial model and $K_{\text{eff}}(x) = O(\log n)$ [3, 8]. That intermediate x can have $K_{\text{eff}}(x) = \Theta(n)$ is witnessed by objects with substantial regularities — a long but genuinely structured pattern — whose optimal model is large. Hence K_{eff} is small at both extremes and large in between. □

5.3 Effective complexity is a genuinely distinct projection

The non-monotonicity has an immediate consequence: effective complexity cannot be recovered from the measures of the other axes.

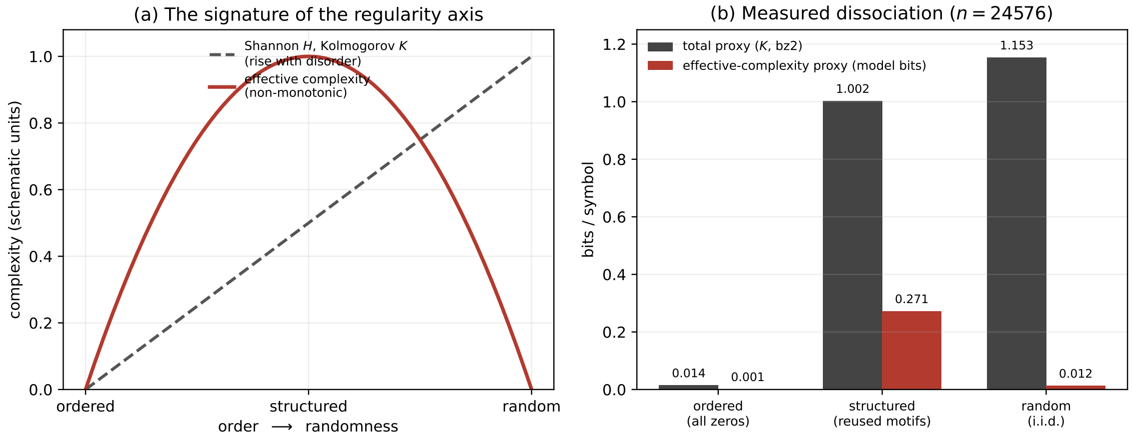

Figure 1: The non-monotonic signature of the regularity axis. **(a)** Schematic: across the order–randomness range, Shannon entropy and Kolmogorov complexity rise monotonically (dashed), while effective complexity is an inverted-U (solid), small at both extremes and large only for the structured middle (after Gell-Mann & Lloyd; structure-function backbone of Vereshchagin & Vitányi). **(b)** A concrete measured dissociation for three n-bit exemplars. A total-complexity proxy (bz2 bits/symbol, grey) rises with disorder; an effective-complexity proxy (the per-symbol bits of reusable structure — the dictionary of recurring length-16 blocks, the model in a two-part code, red) is large only for the structured exemplar. The structured and random exemplars have nearly equal total proxy (1.00 vs. 1.15) yet sharply different effective-complexity proxy (0.27 vs. 0.01): effective complexity is not a function of the description-axis measure alone (Proposition 5.2).

Proposition 5.2 (Dissociation from the description and channel axes)**.** *There exist individuals x, y with $K(x) = K(y) + O(\log n)$ but $K_{\text{eff}}(x) \neq K_{\text{eff}}(y)$ by $\Theta(n)$; likewise there exist sources with equal Shannon entropy but whose typical objects differ in effective complexity. Hence K_{eff} is not a function of the description-axis projection K, nor of the channel-axis projection H, alone: it is a genuinely distinct projection.*

Proof. Take x a structured object with substantial model, $K_{\text{eff}}(x) = \Theta(n)$, and total $K(x) = n + O(\log n)$ in the sense that its description is barely compressible overall; take y incompressible-random with $K(y) = n + O(\log n)$ and $K_{\text{eff}}(y) = O(\log n)$ by Proposition 5.1(ii). Then $K(x) = K(y) + O(\log n)$ while $K_{\text{eff}}(x) - K_{\text{eff}}(y) = \Theta(n)$. The measured exemplars of Figure 1(b) realize this: structured and random strings of nearly equal total proxy (1.00 and 1.15 bits/symbol) have effective-complexity proxies of 0.27 and 0.01 bits/symbol. The entropy version is analogous, since equal-entropy sources can have typical objects of differing structure. □

This is the precise content of "no proportionality" in the refinement morphism. The channel and macrostate axes were related by a rescaling, so that one determined the other up to a constant; the description and regularity axes are related by an additive split in which the model-share varies independently of the total, so that the description-axis measure leaves the regularity-axis measure undetermined. Figure 1 shows the dissociation concretely: two objects the description axis cannot tell apart by total complexity, the regularity axis sharply distinguishes.

5.4 Counterarguments and Responses

Objection 1 (the model-cost proxy in Figure 1 is not effective complexity). *Response.* It is a computable lower-bound illustration, and is labeled as such: the per-symbol bits of recurring length-16 blocks is one concrete two-part model whose cost bounds the structural content from below and exhibits the dissociation. It is no more the exact effective complexity than bz2 was the exact Kolmogorov complexity in the description article; it demonstrates the qualitative fact — structured objects carry substantial model bits, random ones do not — that Proposition 5.1 states exactly and uncomputably.

Objection 2 (the non-monotonicity makes effective complexity ill-behaved as a measure). *Response.* It makes it non-monotonic, which is precisely the behavior the measure was designed to have: a measure of structure *must* be small for both the trivial and the random, or it is a measure of disorder, not structure. The non-monotonicity is the signature of the regularity facet, not a pathology; what would be ill-behaved is a measure of structure that rose monotonically with randomness, for it would not measure structure at all.

5.5 Section Result and Implications

We have exhibited the shape of the regularity axis: its native indeterminacy is the machine constant together with the structure/residue cut; its signature is the non-monotonic hump (Proposition 5.1, Figure 1); and effective complexity is a genuinely distinct projection, not a function of the description- or channel-axis measure alone (Proposition 5.2). The implication is that the cut, which the shape required us to fix, is the axis's characteristic indeterminacy, to which we now turn.

6 The Characteristic Feature: Relativity of the Cut

6.1 Contestedness located

Each axis has a characteristic feature flowing from what its reduction requires. The channel axis's was observer-relativity, the dependence of the entropy on the chosen source model; the description axis's was uncomputability, the price of retaining the whole individual. The regularity axis's characteristic feature is the relativity of the cut, and it is the precise content of effective complexity's reputation for being contested.

Proposition 6.1 (Effective complexity is determined only relative to the cut)**.** *Effective complexity $K_{\text{eff}}(x)$ is well-defined only relative to a choice of the structure/residue cut — equivalently, the model class admitted, or the level of the structure function at which the minimal sufficient statistic is read. Different admissible cuts assign x different effective complexities; given the cut, $K_{\text{eff}}(x)$ is determined (up to the machine constant).*

Proof. The minimal sufficient statistic, and hence K_{eff}, is defined as the elbow of the structure function relative to an admitted class of models (Definition 4.2). Enlarging or restricting the class moves the elbow: a feature counted as structure under a permissive class is counted as residue under a restrictive one, changing $K(M)$ and hence $K_{\text{eff}}(x)$. Thus K_{eff} is a function of x and the cut jointly; fixing the cut fixes the function [3, 4, 8]. □

6.2 Why this dissolves the charge of subjectivity

Effective complexity has been criticized as subjective, on the ground that what one observer counts as regularity another counts as noise, so that the measure depends on the observer's interests. The substrate reading grants the observation and reverses its force. The dependence is real, and it is the relativity of Proposition 6.1; but it is not a subjectivity of the *measure*, it is the native indeterminacy of the *axis*, exactly parallel to the channel axis's dependence on the source model. There, the entropy was objective given the reduction and the observer was spent in choosing the reduction [6]; here, the effective complexity is objective given the cut and the observer is spent in choosing the cut. In both cases the "subjectivity" is the choice the reduction requires, made once, before the functional is evaluated; in neither does it infect the functional. To call effective complexity subjective is to mistake the indeterminacy of the axis for an indeterminacy of the measure, just as calling Shannon entropy subjective mistook the choice of source for a subjectivity of H. The cut must be chosen; once chosen, effective complexity is as objective as entropy.

This is the regularity axis's place in the pattern the volume has been building. Three axes, three characteristic features, each the native indeterminacy of a reduction: the channel axis relative to its source model, the description axis uncomputable as the price of the whole individual, the regularity axis relative to its structure/residue cut. The features look like three separate embarrassments — entropy's subjectivity, complexity's uncomputability, effective complexity's contestedness — and are, on the substrate reading, three instances of one fact: a projection requires a choice of reduction, and the choice is where the apparent defect lives.

6.3 Counterarguments and Responses

Objection 1 (if the cut is free, effective complexity measures nothing definite). *Response.* It measures something definite relative to each cut, exactly as entropy measures something definite relative to each source model. "Free" overstates it: the cut is constrained by the purpose for which structure is being measured, and natural purposes fix natural cuts, as a physical coarse-graining fixes a natural macrostate. What the substrate frame denies is only an *absolute*, cut-free effective complexity — the same denial it made of an absolute, source-free entropy and an absolute, axis-free amount of information (Article 2). Relative to a cut, the measure is definite; absent any cut, the question "how structured is x, full stop" is, like "how much information does the universe contain, full stop," ill-posed at the level of the substrate.

Objection 2 (the parallel with entropy's observer-relativity is too neat to be real). *Response.* The parallel is exact in the one respect claimed: each is the dependence of a projection measure on the reduction that defines its axis, with the measure determinate once the reduction is fixed. It is not claimed that the choices are of the same kind — a source model, a universal machine, and a structure cut are different choices — only that each is the native indeterminacy of its axis, and that the apparent subjectivity or contestedness of each measure is that indeterminacy and not a defect of the functional. That three independent measures exhibit the same structural pattern is evidence for the substrate reading, not an artifact of it.

6.4 Section Result and Implications

We have located the characteristic feature of the regularity axis: effective complexity is determined only relative to the structure/residue cut (Proposition 6.1), and this relativity is the axis's native indeterminacy, dissolving the charge of subjectivity into the precise account that the cut is the choice the reduction requires. The implication is that the three axes of Part II share one pattern — a reduction, a functional, and a characteristic indeterminacy — and that effective complexity's contestedness is its instance of it. We close by marking what the axis cannot see.

7 What the Regularity Axis Cannot See

The regularity axis is the complete theory of structure, and within it the two-part code, the structure function, the minimal sufficient statistic, and the total information answer the questions the axis can pose. But like its predecessors it is blind, from within, to what lies across its borders.

It cannot fix its own cut. The level at which structure gives way to residue is set by the purpose for which structure is measured, and that purpose is not in the regularity projection; the axis takes the cut as given and computes relative to it, exactly as the channel axis took the source model as given. Where the cut comes from — a physical coarse-graining, a task, a context — is a question the regularity axis cannot answer from within, and the substrate frame locates the cut as the reduction's required choice without supplying its content. Nor can the axis see, from within, why its measure relates to the description axis by an additive split rather than a rescaling; that the relation is the refinement morphism $K = K_{\text{eff}} + R$ is a fact about the common domain, the substrate, visible only when the description and regularity projections are seen as reductions of one object [5].

This is the same sense in which the substrate frame held the channel and description theories: it assigns the theory of structure its facet — the regularity axis, with its model-retaining reduction, its two-part functional, its structure function — and affirms its completeness there, reserving for the substrate only the relations to the other facets and the source of the cut. We locate the theory's authors with the generosity the volume owes its predecessors. Gell-Mann and Lloyd [3, 4] built effective complexity and the total information, isolating structure from both order and randomness and giving the regularity facet its measure; Bennett [1] built logical depth, the neighboring measure that finds structure in computation time rather than model size; and Vereshchagin and Vitányi [8] and Gács et al. [2] gave the structure function and the minimal sufficient statistic the rigorous form on which the modern account rests. Each saw the regularity facet truly, and built the theory that the substrate frame now places beside the channel and description theories as the third complete account of a projection.

7.1 Counterarguments and Responses

Objection 1 (if the regularity axis cannot fix its own cut, it is parasitic on the other axes). *Response.* It is built on the description axis — that is the content of the refinement morphism — but it is not parasitic: it isolates a quantity, the structural share of a description, that no other axis measures, and the dissociation of Proposition 5.2 shows that quantity to be independent of the description total. Depending on the description axis for its raw material is not the same as being reducible to it; the regularity axis adds the structure/residue split, which is genuinely its own.

Objection 2 (locating logical depth "beside" effective complexity dodges the question of which is right). *Response.* The question "which is right" presupposes a single correct measure of structure, which the axis no more has than the channel axis had a single correct order of Rényi entropy. Size and time are two resources in which structure can be costly, and effective complexity and logical depth measure the two; which is the relevant structure is fixed by the purpose, like the cut. Locating both on the regularity axis is the answer the substrate frame gives, and it is not a dodge but the same pluralism-within-an-axis the channel article established.

7.2 Section Result and Implications

We have marked what the regularity axis cannot see — the source of its cut and the refinement morphism that relates it to the description axis, both supplied from the substrate — and located Gell-Mann, Lloyd, Bennett, Vereshchagin, and Vitányi as the authors of the complete theory of the structured facet. The implication, carried forward in Part II, is that the macrostate axis of thermodynamic entropy, the last of the principal projections, receives its own located treatment next.

8 Epistemic Status: Proved, Imported, Affirmed

In the manner of the series, we close with an explicit ledger. *Proved (here)* marks results established in this article; *Imported* marks results taken from prior articles, chiefly the refinement morphism of Article 3; *Affirmed (standard)* marks established algorithmic and complexity theory the article endorses and locates without re-deriving.

Claim	**Status**	**Locus**
K_{eff} = regularity-axis projection; $K = K_{\text{eff}} + R$ refinement morphism	Imported	Kriger [5] (Art. 3)
Regularity reduction refines description: model kept, residue discarded	Proved (here)	Def. 3.1, Rem. 3.2
Non-monotonicity of effective complexity (the hump)	Proved (here)	Prop. 5.1
Dissociation: K_{eff} not a function of K or H alone	Proved (here)	Prop. 5.2
Effective complexity determined only relative to the cut	Proved (here)	Prop. 6.1
Contestedness = native indeterminacy of the axis (not subjectivity)	Proved (here)	§6.2
Two-part code; total information	Affirmed (standard)	Gell-Mann and Lloyd [3, 4]
Kolmogorov structure function; minimal sufficient statistic	Affirmed (standard)	Gács et al. [2], Vereshchagin and Vitányi [8]
Logical depth (neighboring regularity-type measure)	Affirmed (standard)	Bennett [1]
Uncomputability of K_{eff} (bears on evaluation, not definition)	Imported	Kriger [7]

The ledger marks the division of labor. The article *imports* the refinement morphism from Article 3 and the uncomputability from Article 6; it *proves* the structural facts about the regularity reduction, the non-monotonic signature, the dissociation, and the relativity of the cut; and it *affirms*, without re-deriving, the established theory of effective complexity, the structure function, and logical depth that it locates.

9 Discussion

The result of this article is a placement, the third of Part II. Effective complexity is located on the regularity axis as the complete theory of structure: the regularity reduction refines the description reduction, splitting the shortest description into a model and a residue and retaining only the model, and the two-part code, the structure function, and the minimal sufficient statistic are the apparatus of that retained model. The placement makes two contributions beyond the bookkeeping. It gives the non-monotonicity of effective complexity its precise statement and a concrete measured demonstration — that structured and random objects of nearly equal total complexity can have sharply different effective complexity — showing the measure to be a genuinely distinct projection rather than a rescaling of complexity. And it locates the contestedness of effective complexity as the relativity of the structure/residue cut, the regularity axis's native indeterminacy, parallel to the channel axis's observer-relativity and the description axis's uncomputability, thereby dissolving the charge of subjectivity into a precise account of what the reduction requires.

The placement completes a pattern. The three axes of Part II so far — channel, description, regularity — each consist of a reduction, a functional, and a characteristic indeterminacy, and the three indeterminacies, which the literature has treated as three separate difficulties, are one fact seen three times: a projection requires a choice of reduction. Entropy's apparent subjectivity, complexity's uncomputability, and effective complexity's contestedness are the channel's source model, the description's universal machine, and the regularity's structure cut. The substrate frame does not remove these choices — they are the price of projecting an axis-free substrate onto a single coordinate — but it locates each as the native indeterminacy of an axis and shows the functional to be determinate once the choice is made. That is the discipline of the classification: to affirm each theory as complete within its axis, to locate the choice its reduction requires, and to reserve for the substrate only the relations among the axes.

9.1 Limitations

Three limitations are material. *First*, the article imports the refinement morphism and the uncomputability from Articles 3 and 6 and does not re-establish them; a reader who doubts those results will find the placement only as secure as its sources. *Second*, the figure's effective-complexity proxy is a computable lower-bound illustration — the bits of recurring fixed-length blocks — and demonstrates the dissociation and the structured-not-random behavior without computing the exact effective complexity, which is uncomputable; the exact non-monotonicity is the content of Proposition 5.1, not of the proxy. *Third*, the relativity of the cut is located but not eliminated: the article does not supply a canonical cut, holding that the cut is fixed by purpose and that an absolute, purpose-free effective complexity is, like an absolute amount of information, ill-posed at the level of the substrate.

9.2 Future Directions

The immediate sequel completes the principal projections of Part II with the macrostate axis of thermodynamic entropy, whose relation to the channel axis is the rescaling rigidity of Article 3 and whose located treatment parallels the three given so far. Beyond the volume, the relativity of the structure/residue cut invites a precise question continuous with Article 3 and with the deformation question raised for the channel and description axes: whether the admissible cuts form a structure — a lattice of model classes, perhaps — whose morphisms could be classified alongside the rescaling, kernel, and refinement morphisms between axes, turning the choice of cut from a free parameter into a position in a classified space. That question connects the present located treatment to the open completeness conjecture of Article 3 and is left to the work that can address it.

Acknowledgments

The author thanks colleagues of the Information Physics Institute and the Institute of Integrative and Interdisciplinary Research for sustained discussion of the foundations of the series.

References

[1] C. H. Bennett. Logical depth and physical complexity. In R. Herken, editor, *The Universal Turing Machine: A Half-Century Survey*, pages 227–257. Oxford University Press, Oxford, 1988.

[2] P. Gács, J. T. Tromp, and P. M. B. Vitányi. Algorithmic statistics. *IEEE Transactions on Information Theory*, 47(6):2443–2463, 2001.

[3] M. Gell-Mann and S. Lloyd. Information measures, effective complexity, and total information. *Complexity*, 2(1):44–52, 1996.

[4] M. Gell-Mann and S. Lloyd. Effective complexity. In M. Gell-Mann and C. Tsallis, editors, *Nonextensive Entropy: Interdisciplinary Applications*, pages 387–398. Oxford University Press, Oxford, 2004.

[5] B. Kriger. The projection theorem: how known information measures descend from the substrate. The Information Substrate Theory (IST), Volume VII, Article 3. Zenodo, 2026.

[6] B. Kriger. The axis of the channel: Shannon information as a substrate projection. The Information Substrate Theory (IST), Volume VII, Article 5. Zenodo, 2026.

[7] B. Kriger. The axis of description: Kolmogorov complexity and the algorithmic projection. The Information Substrate Theory (IST), Volume VII, Article 6. Zenodo, 2026.

[8] N. K. Vereshchagin and P. M. B. Vitányi. Kolmogorov's structure functions and model selection. *IEEE Transactions on Information Theory*, 50(12):3265–3290, 2004.

A Code for Figure 1

The following Python script generates Figure 1. Panel (a) plots the schematic curves; panel (b) measures, for three n-bit exemplars, a total-complexity proxy (bz2) and an effective-complexity proxy (the per-symbol bits of recurring length-L blocks, a computable two-part model cost).

```
import bz2
from collections import Counter
import numpy as np
rng = np.random.default_rng(7); n = 24576; L = 16

def packed(b):
    pad = (-len(b)) % 8
    if pad: b = np.concatenate([b, np.zeros(pad, np.uint8)])
    return np.packbits(b).tobytes()

def k_proxy(b):                       # total-complexity proxy (bits/symbol)
    return 8.0 * len(bz2.compress(packed(b), 9)) / len(b)

def eff_proxy(b, L=L):                # effective-complexity proxy (model bits/sym)
    nb = len(b) // L
    blocks = [bytes(b[i*L:(i+1)*L].tolist()) for i in range(nb)]
    reused = [k for k, v in Counter(blocks).items() if v >= 2]
    return (len(reused) * L) / len(b)

ordered    = np.zeros(n, np.uint8)                       # K~0.014, Keff~0.001
D, r = 512, 3                                            # dictionary of motifs
motifs = (rng.random((D, L)) < 0.5).astype(np.uint8)
idx = rng.integers(0, D, size=n // L)
structured = np.concatenate([motifs[i] for i in idx])[:n]   # K~1.002, Keff~0.271
random_    = (rng.random(n) < 0.5).astype(np.uint8)         # K~1.153, Keff~0.012
```

The Axis of the Macrostate: Thermodynamic Entropy as a Substrate Projection

The complete theory of the coarse-grained facet, located: microstate counting, the rescaling to the channel, and the second law

Boris Kriger[1,2]

[1] Information Physics Institute
boris.kriger@informationphysicsinstitute.net

[2] Institute of Integrative and Interdisciplinary Research, Toronto
boriskriger@interdisciplinary-institute.org

ORCID: https://orcid.org/0009-0001-0034-2903

The Information Substrate Theory (IST) — Volume VII, Article 8
Part II: The Projections — classification of information theories

Abstract

Thermodynamic entropy is the oldest of the information measures and the only one with units. This article, completing the principal projections of the classificatory part of the volume, locates Boltzmann–Gibbs entropy on the macrostate axis of the substrate. The macrostate reduction retains a coarse-grained description — a count of the microstates compatible with a macroscopic specification — and discards the individual microstate the system actually occupies. Boltzmann's $S = k_B \ln W$ is the special case, at equiprobable microstates, of Gibbs's $S = -k_B \sum_i p_i \ln p_i$, and both are k_B times the Shannon entropy of the microstate distribution: the macrostate axis is the channel axis equipped with a physical coarse-graining and a dimensional unit. This is the rescaling rigidity of Article 3, $S = k_B H$, and we develop it from the macrostate side, noting the feature that distinguishes this axis from the others: its unit is not free. Where the channel axis left the logarithm base to choice and the description axis left an additive machine constant, the macrostate axis carries the Boltzmann constant $k_B = 1.380649 \times 10^{-23}$ J/K, fixed exactly by the 2019 redefinition of the SI; the rescaling constant is pinned by physics, not chosen. What the macrostate axis does leave free is the coarse-graining, its native indeterminacy, parallel to the channel's source model, the description's machine, and the regularity's cut. The second law is then located as the lossiness of that coarse-graining: under reversible, measure-preserving micro-dynamics the fine-grained entropy is conserved by Liouville's theorem, yet the coarse-grained entropy increases to equilibrium, an increase produced by the coarse-graining and not by the dynamics. We demonstrate this with a reversible mixing map and locate the thermodynamic arrow of time as the macrostate-axis face of the lossiness of projection. Boltzmann, Gibbs, and Jaynes are located as the authors of the complete theory of the coarse-grained facet.

Keywords: thermodynamic entropy; Boltzmann–Gibbs entropy; macrostate axis; coarse-graining; second law; arrow of time; Boltzmann constant; Liouville's theorem; foundations of physics.

Contents

1 Introduction

The three preceding articles of Part II located the channel, description, and regularity axes [7–9]. The measures they placed — Shannon entropy, Kolmogorov complexity, effective complexity — share one striking feature: they are all dimensionless. They count bits, or program lengths, or model sizes, but they carry no units of physics. Thermodynamic entropy is different. It was the first quantity to be called entropy, it predates information theory by most of a century, and it carries units of energy over temperature: joules per kelvin. A classification of the information measures that did not reach thermodynamic entropy would have left out the one with the strongest claim to physical reality, and the one whose relation to the others has caused the most confusion and the most insight. This article completes the principal projections by locating Boltzmann–Gibbs entropy on the macrostate axis of the substrate.

The placement turns on a relation already proved. The Projection Theorem of Article 3 [6] identified the macrostate–channel relation as a *rescaling*: thermodynamic entropy is the Boltzmann constant times an information entropy, $S = k_B H$. The macrostate reduction takes a physical system, fixes a macroscopic description, and counts the microstates compatible with it; the Shannon entropy of the resulting microstate distribution, multiplied by k_B, is the thermodynamic entropy. Boltzmann's celebrated $S = k_B \ln W$ is the equiprobable case of Gibbs's $S = -k_B \sum_i p_i \ln p_i$ [3, 4], and both are k_B times the channel-axis functional applied to the microstate distribution. The macrostate axis is, in this precise sense, the channel axis with two additions: a physical coarse-graining that fixes which configurations count as the same macrostate, and a dimensional unit that fixes the scale.

Two features distinguish this axis from its predecessors, and they organize the article. The first concerns the unit. On the channel axis the unit — the base of the logarithm — was free, a matter of bits or nats; on the description axis the analogue was an additive machine constant. On the macrostate axis the unit is not free. It is the Boltzmann constant $k_B = 1.380649 \times 10^{-23}$ J/K, fixed exactly by the 2019 redefinition of the International System of Units [2]. The rescaling constant of $S = k_B H$ is pinned by physics, not chosen, and this is the macrostate axis's distinctive structural fact: it is the channel axis whose free unit has been determined by the world. The second feature concerns what the macrostate axis does leave free, which is the coarse-graining — the choice of which macrovariables, which partition of the microstate space. This is the macrostate axis's native indeterminacy, parallel to the channel's source model, the description's universal machine, and the regularity's structure cut, and it is the seat of the axis's characteristic phenomenon.

That phenomenon is the second law. The non-decrease of thermodynamic entropy, and with it the thermodynamic arrow of time, is on the substrate reading the lossiness of the coarse-graining made manifest. Under reversible, measure-preserving micro-dynamics the fine-grained entropy of the exact phase-space density is conserved — Liouville's theorem — yet the coarse-grained entropy increases to equilibrium, because the reversible dynamics stretches and folds the density into filaments that the coarse-graining can no longer resolve. The increase is produced by the coarse-graining, not by the dynamics; the second law is the macrostate-axis face of the theme that runs through the whole volume, that projection loses. We demonstrate this with a reversible mixing map and locate the thermodynamic arrow accordingly, deferring its full development to Part III.

We close, as the companion articles did, by marking what the macrostate axis cannot see and locating its authors generously. Boltzmann, Gibbs, and Jaynes built the complete theory of the coarse-grained facet; the substrate frame assigns it its axis and explains why that axis relates to the channel axis by a rescaling whose constant the world has fixed.

2 Logical Structure of the Argument

The argument completes the pattern of Part II for the fourth and most physical axis.

Step 1 (Section 3): The macrostate reduction.
The macrostate reduction retains a coarse-grained description and the count of compatible microstates, and discards the individual microstate (Definition 3.1).

Step 2 (Section 4): The apparatus.
Boltzmann's $S = k_{\mathrm{B}} \ln W$ is the equiprobable case of Gibbs's entropy (Proposition 4.1); we exhibit temperature, the fundamental relation, and the second law as macrostate-axis constructs.

Step 3 (Section 5): The shape of the axis.
The native unit is the Boltzmann constant, fixed by physics rather than chosen (Proposition 5.1); the macrostate–channel relation is the rescaling rigidity $S = k_{\mathrm{B}} H$, developed from the macrostate side.

Step 4 (Section 6): Coarse-graining and the second law.
The coarse-graining is the axis's native indeterminacy (Proposition 6.1); the second law is the lossiness of the coarse-graining under reversible micro-dynamics (Proposition 6.2, Figure 1), and the thermodynamic arrow is the macrostate-axis face of the lossiness of projection.

Step 5 (Section 7): What the macrostate axis cannot see.
The axis is the theory of the coarse-grained facet and is blind, from within, to the source of its coarse-graining and to why its constant is what it is; we locate Boltzmann, Gibbs, and Jaynes, and forward Maxwell's demon to the cost axis.

Step 6 (Section 8): Epistemic status.
An explicit ledger separates what is proved here, what is imported, and what is affirmed as standard physics.

The dependencies are linear: Steps 2–4 operate on the reduction of Step 1; Step 5 sets the whole within the cross-axis frame of Article 3. The thesis advanced is that

Boltzmann–Gibbs entropy is the complete and correct theory of the coarse-grained facet, that its unit is the channel unit pinned by physics, and that the second law is its coarse-graining's loss.

3 The Macrostate Reduction

3.1 Macrostate retained, microstate discarded

We begin from the reduction that defines the axis.

Definition 3.1 (Macrostate reduction)**.** The *macrostate reduction* ρ_{ma} fixes a coarse-graining of a physical system's microstate space — a partition into macrostates, each a set of microstates indistinguishable at the chosen macroscopic level — and sends the system to its macrostate together with the distribution over the microstates that macrostate contains. It retains the macrostate and the microstate count (or distribution); it discards which individual microstate the system actually occupies. The measure on the macrostate projection is the Boltzmann–Gibbs entropy.

What the macrostate reduction discards is the individual microstate: the exact positions and momenta of every molecule, the precise configuration within the macrostate. What it retains is the macrostate — the equivalence class — and how many microstates realize it, or with what weights. This places the macrostate axis close to the channel axis but not identical to it. The channel axis retained the statistics of repetition of distinguished outcomes; the macrostate axis retains the statistics of microstates within a physically fixed coarse-graining. The two differ in what fixes the outcomes: on the channel axis a source model, on the macrostate axis a physical partition of phase space. The macrostate axis is the channel axis with the outcomes supplied by physics and the unit supplied by physics, as Section 5 makes precise.

Remark 3.2 (Macrostate and microstate)**.** The distinction between macrostate and microstate is the whole content of the reduction. A microstate is a complete specification of the system — the description-axis individual of Article 6, in the physical setting. A macrostate is a coarse-graining of microstates, an equivalence class under a chosen macroscopic description. The macrostate reduction is therefore a coarse-graining of the microstate space, and thermodynamic entropy measures the size of the equivalence class, not the individual within it. This is why thermodynamic entropy is a property of the macrostate and the coarse-graining, never of the substrate (Article 2): the substrate has no macrostate until a coarse-graining is fixed.

3.2 Counterarguments and Responses

Objection 1 (the macrostate reduction is just the channel reduction, so the macrostate axis is not distinct). *Response.* It is the channel reduction with two physical additions, and the additions make the difference. The coarse-graining fixes the outcomes by physics rather than by a free source model, and the unit fixes the scale by physics rather than by a free logarithm base. The resulting measure carries units, enters the first and second laws, and relates to energy and temperature — none of which the bare channel entropy does. The rescaling $S = k_{\mathrm{B}}H$ relates the two axes precisely (Section 5); it does not identify them.

Objection 2 (discarding the microstate is the same discarding the description axis already performed). *Response.* No: the description axis *retained* the individual microstate — that was its defining choice (Article 6). The macrostate axis discards it, retaining only the coarse-grained class. The two axes make opposite choices about the individual, which is why the microstate is the description-axis object and the macrostate is its coarse-graining (Remark 3.2). Thermodynamic entropy is large exactly when the description axis would need many bits to single out the microstate within the macrostate.

3.3 Section Result and Implications

We have fixed the macrostate reduction as the physical coarse-graining that retains the macrostate and the microstate count and discards the individual microstate (Definition 3.1, Remark 3.2). The implication is that thermodynamic entropy is a property of a coarse-graining, close to the channel entropy but supplied by physics with both its outcomes and its unit — a profile the Boltzmann–Gibbs apparatus exhibits exactly, as we now show.

4 The Apparatus

4.1 Boltzmann, Gibbs, and the microstate distribution

The functional native to the macrostate axis has two classical forms, and their relation is the first thing to fix.

Proposition 4.1 (Boltzmann is Gibbs at equiprobability)**.** *Let a macrostate contain W microstates with probabilities $p_1, \ldots, p_W$. The Gibbs entropy is*

$$S = -k_{\mathrm{B}} \sum_{i=1}^{W} p_i \ln p_i. \tag{1}$$

When the microstates are equiprobable, $p_i = 1/W$, this reduces to the Boltzmann entropy

$$S = k_{\mathrm{B}} \ln W. \tag{2}$$

In both cases $S = k_{\mathrm{B}}H$, where H is the Shannon entropy (in nats) of the microstate distribution [3–5].

Proof. Substituting $p_i = 1/W$ into (1) gives $-k_\mathrm{B}\sum_i (1/W)\ln(1/W) = k_\mathrm{B}\ln W$, which is (2). That (1) equals k_B times the Shannon entropy of $(p_1, \ldots, p_W)$ is immediate from the definition of $H = -\sum_i p_i \ln p_i$ [5]. □

The proposition fixes the relation that the rest of the article develops: thermodynamic entropy, in either the Boltzmann or the Gibbs form, is k_B times the channel-axis functional applied to the microstate distribution of a macrostate. Boltzmann's form counts microstates and is exact for the equiprobable case that maximum entropy selects at fixed energy; Gibbs's form weights them and covers the general ensemble. Both live on the macrostate axis, and both differ from the channel entropy only by the physical coarse-graining that supplies the microstates and the constant k_B that supplies the unit.

4.2 Temperature, the fundamental relation, and the second law

On the macrostate axis the entropy is not an isolated number but the hinge of thermodynamics. Temperature is defined through it, $1/T = (\partial S/\partial U)_{V,N}$, so that the entropy's sensitivity to energy is the inverse temperature; the fundamental relation $dU = T\,dS - p\,dV + \mu\,dN$ makes the entropy the potential conjugate to temperature, through which heat and work are distinguished [11]. And the *second law* states that for an isolated system the thermodynamic entropy does not decrease, $dS \geq 0$, reaching its maximum at equilibrium. These are the operational content of the macrostate axis, the analogue of the coding theorems on the channel axis: they are what make the macrostate entropy a theory of heat and not merely a functional. On the substrate reading they are constructs on the coarse-grained projection, and the second law in particular is the axis's characteristic phenomenon, treated in Section 6.

4.3 Counterarguments and Responses

Objection 1 (the Boltzmann and Gibbs entropies are known to disagree in some cases, so calling them one functional is wrong). *Response.* They disagree precisely where the microstate distribution is non-uniform and the system is out of equilibrium, and Proposition 4.1 marks this: Boltzmann's form is the equiprobable case, Gibbs's the general one, and both are $k_\mathrm{B}H$ of the relevant microstate distribution. The substrate reading does not paper over the distinction; it locates it as the difference between the uniform and the weighted microstate distribution fed to the channel functional, which is exactly the difference the debate concerns.

Objection 2 (temperature and the laws are physics, not information, so locating them on an "axis" overreaches). *Response.* The location does not reduce them to information; it observes that the entropy through which they are defined is k_B times a channel-axis functional of a physical coarse-graining. Temperature, heat, and work remain physical; what the substrate reading adds is that the entropy on which they hinge is the macrostate projection, related to the channel projection by the rescaling $S = k_B H$. The physics is untouched; its entropy is placed.

4.4 Section Result and Implications

We have exhibited the apparatus of the macrostate axis — the Boltzmann and Gibbs entropies as the equiprobable and general forms of $k_B H$ (Proposition 4.1), and temperature, the fundamental relation, and the second law as the constructs that make it a theory of heat. The implication is that the macrostate axis carries a complete physical theory, and that its entropy is the channel functional supplied by physics with outcomes and a unit. We turn to the shape of the axis, where the determined unit lives.

5 The Shape of the Axis

5.1 The unit is not free

Every axis of Part II has carried a native indeterminacy in its unit or its reduction. The macrostate axis is the one whose *unit* is not free.

Proposition 5.1 (The rescaling rigidity with a determined constant)**.** *The macrostate–channel relation is the rescaling*

$$S = k_B H, \tag{3}$$

where H is the channel-axis (Shannon) entropy of the microstate distribution and k_B is the Boltzmann constant. Because the channel axis is exactly additive (Article 5) and the macrostate reduction feeds a physical microstate distribution to the channel functional, the relation is a proportionality, not merely a correlation; and the constant of proportionality is not free but fixed, $k_B = 1.380649 \times 10^{-23}$ J/K exactly, by the 2019 redefinition of the SI [2, 6].

Proof. That $S = k_B H$ is Proposition 4.1. That the relation is a rescaling — a proportionality with constant k_B — rather than a looser dependence is the rescaling rigidity of Article 3 [6]: exact additivity of the channel functional under independent composition forces any extensive measure built on the same microstate distribution to be proportional to it, and the proportionality constant is the unit ratio. Since 2019 the SI fixes k_B exactly, so the unit ratio is a determined physical constant, not a convention [2]. □

This is the macrostate axis's distinctive structural fact. On the channel axis the unit was a free choice of logarithm base; the same information entropy could be read in bits or nats, and nothing in the world fixed the choice. On the macrostate axis the corresponding unit is the Boltzmann constant, and the world has fixed it: a kelvin and a joule are defined so that k_B takes a particular exact value. The rescaling $S = k_B H$ therefore has a determined constant, and thermodynamic entropy is the channel entropy of a microstate distribution read in the units physics has pinned. The contrast sharpens the placement: the macrostate axis is the channel axis with its free unit removed by the world and its outcomes supplied by a physical coarse-graining.

5.2 Empirical anchoring

The determined unit is not a stipulation but a measured and now fixed fact, and two anchors make this concrete. First, absolute entropies computed from microstate counting — the Sackur–Tetrode entropy of a monatomic ideal gas, for instance — agree with calorimetric measurements, confirming that $S = k_B \ln W$ with the physical k_B describes real systems and not merely a formal analogy [11]. Second, the Landauer bound on the energy of erasing one bit, $k_B T \ln 2$, has been measured directly in a single-bit memory [1, 10], and its value is set by the same k_B; the rescaling constant that links the channel and macrostate axes is the same constant that prices a bit's erasure at the cost axis. That measured bridge is the entry to the cost axis, which the next article of Part II treats. The macrostate axis's unit is thus anchored twice over: in the absolute entropies of matter and in the measured energetics of information erasure.

5.3 Counterarguments and Responses

Objection 1 (fixing k_B by the SI is a human convention, so the unit is free after all). *Response.* The numerical value of k_B reflects the chosen sizes of the joule and the kelvin, but the *physical* content — that thermodynamic entropy and information entropy are related by a universal constant with dimensions of energy over temperature — is not conventional, and the constant's universality across all systems is a fact about the world, not a choice. What the channel axis lacked and the macrostate axis has is exactly this: a single dimensional constant, the same for every system, relating the entropy to energy and temperature. That the SI now fixes its numerical value exactly removes the last measurement uncertainty; it does not make the relation conventional.

Objection 2 (the rescaling rigidity assumes extensivity, which fails for long-range or gravitating systems). *Response.* It assumes the additivity of the channel functional under independent composition, which is exact for the channel axis (Article 5) and underwrites extensive thermodynamic entropy for short-range systems. Where extensivity fails — long-range interactions, gravitation — the appropriate entropy may be a non-additive (for instance Tsallis) form, which Article 5 located as a deformed channel axis. The substrate reading accommodates this without strain: non-extensive thermodynamics is the macrostate axis built on a deformed channel functional, and the rescaling becomes the deformed-axis relation rather than the exactly additive one. The rigidity holds for the additive case and deforms with the functional.

5.4 Section Result and Implications

We have exhibited the shape of the macrostate axis: its unit is the Boltzmann constant, fixed by physics rather than chosen, and the macrostate–channel relation is the rescaling rigidity $S = k_{\mathrm{B}}H$ with a determined constant (Proposition 5.1), anchored in absolute entropies and the measured Landauer bound. The implication is that the macrostate axis is the channel axis with a world-fixed unit and physical outcomes. What it does leave free is the coarse-graining, to which — and to the second law it produces — we now turn.

6 Coarse-Graining and the Second Law

6.1 The coarse-graining is the native indeterminacy

Where the unit of the macrostate axis is fixed, its reduction is not. The choice of coarse-graining — which macrovariables, which partition of the microstate space into macrostates — is the macrostate axis's native indeterminacy, the analogue of the channel's source model, the description's machine, and the regularity's cut.

Proposition 6.1 (Coarse-graining relativity)**.** *Thermodynamic entropy is determined only relative to a choice of coarse-graining: the partition of the microstate space into macrostates. Different admissible coarse-grainings assign a system different entropies, since they count different numbers of microstates as compatible with "the same" macrostate. Given the coarse-graining, the entropy is determined (up to the fixed unit* k_{B}*).*

Proof. The Boltzmann–Gibbs entropy (1) is a functional of the microstate distribution *within a macrostate*, and the macrostate is defined by the coarse-graining (Definition 3.1). A finer coarse-graining resolves more, placing fewer microstates in each macrostate and lowering W; a coarser one raises it. Thus S is a function of the system and the coarse-graining jointly; fixing the coarse-graining fixes the function [5]. □

This is the same structural fact met three times before, in physical dress. The apparent subjectivity of thermodynamic entropy — the observation, due to Jaynes, that "the same" system has different entropies under different descriptions [5] — is not a subjectivity of the functional but the native indeterminacy of the axis. Given a coarse-graining, the entropy is as objective as energy; the coarse-graining is the choice the reduction requires, made once, before the functional is evaluated. The macrostate axis thus completes the pattern of Part II: a reduction, a functional, and a characteristic indeterminacy that the literature has read as subjectivity but that the substrate frame locates as the choice projecting an axis-free substrate onto a single coordinate requires.

6.2 The second law as the lossiness of the coarse-graining

The coarse-graining does more than relativize the entropy; under dynamics it produces the second law.

Proposition 6.2 (The second law as coarse-graining loss)**.** *Let an isolated system evolve under reversible, measure-preserving micro-dynamics. By Liouville's theorem the fine-grained entropy of the exact phase-space density is conserved. Yet the coarse-grained entropy — the Boltzmann–Gibbs entropy of the density averaged over the cells of a fixed coarse-graining — is non-decreasing and rises to its equilibrium maximum, because the reversible dynamics stretches and folds the density into filaments finer than the coarse-graining can resolve. The entropy increase is produced by the coarse-graining, not by the micro-dynamics.*

Proof sketch. Liouville's theorem gives that a measure-preserving flow leaves the fine-grained Gibbs entropy $-k_{\mathrm{B}}\int \rho \ln \rho$ invariant, since ρ is transported without change of its values. The coarse-grained density $\bar{\rho}$, obtained by averaging ρ over each cell, has entropy $-k_{\mathrm{B}}\int \bar{\rho} \ln \bar{\rho} \geq -k_{\mathrm{B}}\int \rho \ln \rho$ by the concavity of $-\rho \ln \rho$ (Jensen's inequality), with equality only when ρ is constant on cells. A mixing flow drives ρ to filament across cells, so $\bar{\rho}$ approaches uniformity and its entropy rises to the maximum $k_{\mathrm{B}} \ln(\text{cells})$, while the fine-grained entropy stays put [11]. The gap between the two is the information the coarse-graining cannot resolve; the second law is its growth. □

The figure makes the proposition concrete. The reversible cat map loses no information — it is a bijection, and its inverse would unmix the blob exactly — yet the coarse-grained entropy climbs to equilibrium, because the coarse-graining cannot follow the filaments the map creates. Run the map backward and the same rise occurs, which is the puzzle of the arrow: the micro-dynamics is symmetric in time, and the asymmetry enters with the coarse-graining and the low-entropy initial macrostate, not with the dynamics. On the substrate reading the thermodynamic arrow of time is therefore the macrostate-axis instance of the asymmetry the volume develops elsewhere: physical projection loses, and the second law is the precise statement of the loss on the macrostate axis. The full development of the arrow — why the initial macrostate was low-entropy, how the coarse-graining's loss relates to the data-processing inequality on the channel axis [7] and to the uncomputability of inversion on the description axis [8] — belongs to Part III; here we mark that the second law is the coarse-graining's loss and not the micro-dynamics'.

6.3 Counterarguments and Responses

Objection 1 (the second law is a fact about physics, not about a coarse-graining we choose). *Response.* The non-decrease of *coarse-grained* entropy is a fact about physics *and* a coarse-graining: the fine-grained entropy does not increase at all (Liouville), so the increase that the second law reports is the increase of the

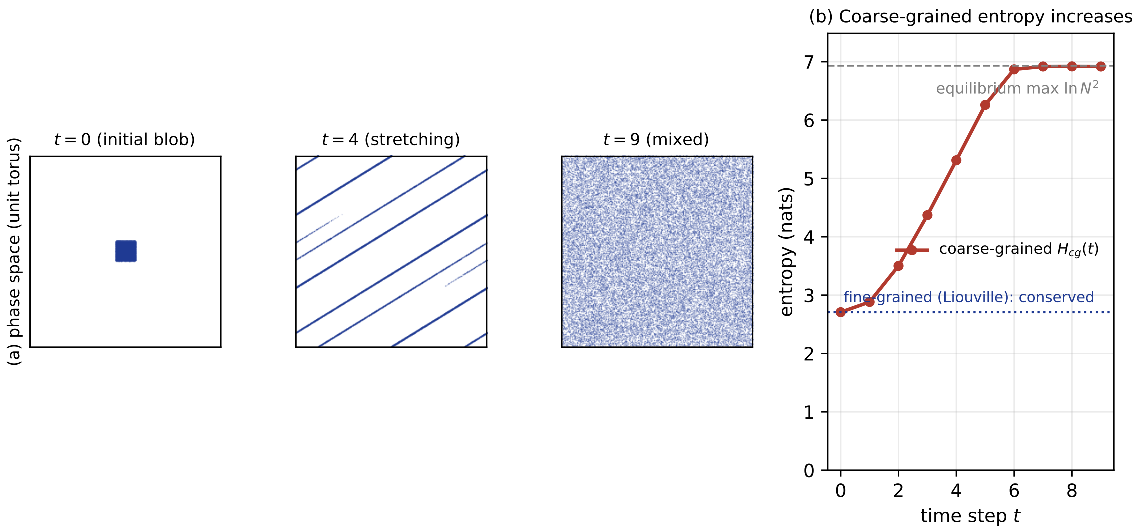

Figure 1: The second law as the lossiness of the coarse-graining. An ensemble of points evolves under Arnold's cat map on the unit torus, $(x, y) \mapsto (2x + y,\, x + y) \bmod 1$, a deterministic, reversible, area-preserving map. **(a)** Phase-space snapshots: a compact initial blob is stretched and folded (mixed) across the torus. **(b)** The coarse-grained entropy $H_{cg}(t)$ of the cell-occupancy distribution on a 32×32 grid rises monotonically from 2.71 to 6.92 nats, reaching the equilibrium maximum $\ln(32^2) = 6.93$ nats. The fine-grained (Liouville) entropy is conserved by the measure-preserving dynamics (blue reference). The increase is produced entirely by the coarse-graining: the reversible micro-dynamics loses nothing, but the filaments it creates fall below the grid's resolution. This is the macrostate-axis face of the lossiness of projection.

coarse-grained entropy, which presupposes a coarse-graining. This is not a denial of the second law but its precise location: the law holds for the coarse-grained entropy, the only entropy that increases, and the coarse-graining is the macrostate axis's reduction. The physics — mixing, equilibration, the approach to the maximum — is untouched; what is located is the sense in which the law concerns a projection.

Objection 2 (if running the map backward also increases coarse-grained entropy, the arrow is not explained). *Response.* It is located, not yet explained, and the article says so. The time-symmetry of the micro-dynamics means the coarse-graining alone does not pick a direction; the direction comes from the low-entropy initial macrostate, the boundary condition whose origin is a cosmological question reserved for Part III. What the present article establishes is the weaker and secure claim that the entropy increase is the coarse-graining's loss rather than the dynamics' — which is what locates the arrow on the macrostate axis and connects it to the lossiness of projection. The origin of the low-entropy past is a separate question, flagged and deferred.

6.4 Section Result and Implications

We have located the macrostate axis's characteristic feature: the coarse-graining is its native indeterminacy (Proposition 6.1), and the second law is the lossiness of that coarse-graining under reversible micro-dynamics (Proposition 6.2, Figure 1), with the thermodynamic arrow the macrostate-axis face of the lossiness of projection. The implication is that the four axes of Part II share one pattern — a reduction, a functional, a characteristic indeterminacy — and that the second law is the macrostate axis's instance of the volume's central asymmetry. We close by marking what the axis cannot see.

7 What the Macrostate Axis Cannot See

The macrostate axis is the complete theory of the coarse-grained facet, and within it the Boltzmann–Gibbs entropy, temperature, the fundamental relation, and the second law answer the questions the axis can pose. But like its predecessors it is blind, from within, to what lies across its borders.

It cannot fix its own coarse-graining. Which macrovariables are the right ones, which partition of phase space counts as "the same" macrostate, is set by the physical situation and the measurements available, and that is not in the macrostate projection; the axis takes the coarse-graining as given and computes relative to it, as the channel axis took the source model as given. Nor can it see, from within, why its constant is what it is: that the rescaling $S = k_B H$ holds with k_B universal across all systems is a fact about the common domain — that thermodynamic and information entropy are reductions of one substrate along the macrostate and channel axes — visible only from the substrate [6]. And it cannot, from within, resolve Maxwell's demon. The demon, who would lower a gas's entropy by sorting its molecules using information about their microstates, sets the macrostate axis against the cost axis: the entropy the demon removes from the gas is paid for in the thermodynamic cost of acquiring and erasing the information, $k_B T \ln 2$ per bit [1, 10]. That resolution lives at the junction of the macrostate and cost axes, and the cost axis is the subject of the next article of Part II; the macrostate axis records the demon's threat but cannot, alone, dissolve it.

This is the same sense in which the substrate frame held the channel, description, and regularity theories: it assigns the theory of thermodynamic entropy its facet — the macrostate axis, with its coarse-grained reduction, its $k_B H$ functional, its laws — and affirms its completeness there, reserving for the substrate the source of the coarse-graining, the universality of the constant, and the relations to the other axes. We locate the theory's authors with the generosity the volume owes its predecessors. Boltzmann [3] gave entropy its microstate-counting meaning, $S = k_B \ln W$, and tied the second law to the approach to the most probable macrostate; Gibbs [4] gave the ensemble form that covers the general microstate distribution; and Jaynes [5] made explicit that thermodynamic entropy is the Shannon entropy of the macrostate-conditioned microstate distribution, the bridge that the substrate frame reads as the rescaling rigidity. Each saw the coarse-grained facet truly, and built the theory that the substrate frame now places beside the channel, description, and regularity theories as the fourth complete account of a projection.

7.1 Counterarguments and Responses

Objection 1 (if the macrostate axis cannot fix its coarse-graining or its constant, it is incomplete). *Response.* It is complete *within* its axis — given a coarse-graining and the constant, every thermodynamic question has its answer — and incomplete only in the sense every projection is: it cannot supply, from within, the choice its reduction requires or the cross-axis relations the substrate induces. That is not a defect of thermodynamics but the mark of a projection, shared by the channel, description, and regularity axes. The substrate frame supplies the coarse-graining's source and the constant's universality from outside the axis, where they live.

Objection 2 (Maxwell's demon has been resolved by Landauer's principle already, so deferring it to the cost axis is evasive). *Response.* The resolution *is* Landauer's principle, and Landauer's principle is the cost axis: the demon is dissolved exactly by recognizing that the information it uses carries a thermodynamic erasure cost $k_B T \ln 2$, which is a cost-axis quantity. Deferring the demon to the cost axis is therefore not evasion but correct placement: the resolution belongs to the junction of the macrostate and cost axes, and the cost axis is treated next. The macrostate axis alone sees the demon's threat to the second law; it takes the cost axis to see the payment that answers it.

7.2 Section Result and Implications

We have marked what the macrostate axis cannot see — the source of its coarse-graining, the universality of its constant, and the resolution of Maxwell's demon, the last at the junction with the cost axis — and located Boltzmann, Gibbs, and Jaynes as the authors of the complete theory of the coarse-grained facet. The implication, carried forward in Part II, is that the cost axis, where information acquires a thermodynamic price and the demon is paid off, receives its treatment next.

8 Epistemic Status: Proved, Imported, Affirmed

In the manner of the series, we close with an explicit ledger. *Proved (here)* marks results established in this article; *Imported* marks results taken from prior articles, chiefly the rescaling rigidity of Article 3; *Affirmed (physics)* marks established physics and metrology the article endorses and locates without re-deriving.

Claim	**Status**	**Locus**
$S = k_{\mathrm{B}}H$ rescaling rigidity (macrostate–channel)	Imported	Kriger [6] (Art. 3)
Exact additivity of the channel functional	Imported	Kriger [7] (Art. 5)
Macrostate reduction: macrostate kept, microstate discarded	Proved (here)	Def. 3.1
Boltzmann is Gibbs at equiprobability; both $= k_{\mathrm{B}}H$	Proved (here)	Prop. 4.1
Rescaling with a determined constant (k_{B} fixed by physics)	Proved (here)	Prop. 5.1
Coarse-graining relativity (native indeterminacy)	Proved (here)	Prop. 6.1
Second law as lossiness of coarse-graining	Proved (here)	Prop. 6.2
Boltzmann $S = k_{\mathrm{B}} \ln W$; Gibbs ensemble entropy	Affirmed (physics)	Boltzmann [3], Gibbs [4]
Entropy as inference; coarse-graining dependence	Affirmed (physics)	Jaynes [5]
Temperature, fundamental relation, second law	Affirmed (physics)	Pathria and Beale [11]
$k_{\mathrm{B}} = 1.380649 \times 10^{-23}\,\mathrm{J/K}$ exact (SI 2019)	Affirmed (metrology)	BIPM [2]
Landauer bound $k_{\mathrm{B}}T \ln 2$, measured	Affirmed (physics)	Bérut et al. [1], Landauer [10]
Maxwell's demon resolution (cost axis)	Deferred	Article 9
Origin of the low-entropy past (arrow)	Deferred	Part III

The ledger marks the division of labor. The article *imports* the rescaling rigidity and the additivity of the channel functional from Articles 3 and 5; it *proves* the structural facts about the macrostate reduction, the Boltzmann–Gibbs relation, the determined constant, the coarse-graining relativity, and the second law as coarse-graining loss; it *affirms* the established physics and metrology it locates; and it *defers* the resolution of Maxwell's demon to the cost axis and the origin of the low-entropy past to Part III.

9 Discussion

The result of this article is a placement, the fourth and last of the principal projections of Part II. Thermodynamic entropy is located on the macrostate axis as the complete theory of the coarse-grained facet: the macrostate reduction fixes a physical coarse-graining and counts the microstates compatible with it, the Boltzmann–Gibbs entropy is k_{B} times the channel functional of the resulting microstate distribution, and temperature, the fundamental relation, and the second law are the apparatus that makes it a theory of heat. The placement makes two contributions beyond the bookkeeping. It identifies the macrostate axis as the channel axis with its unit determined by physics: where the channel's logarithm base was free, the macrostate's constant is the Boltzmann constant, fixed exactly by the SI, so the rescaling $S = k_B H$ carries a world-fixed proportionality. And it locates the second law as the lossiness of the coarse-graining — the fine-grained entropy conserved by Liouville's theorem, the coarse-grained entropy rising because the reversible dynamics filaments below the coarse-graining's resolution — placing the thermodynamic arrow of time as the macrostate-axis face of the lossiness of projection.

The placement completes the pattern of Part II. Four axes — channel, description, regularity, macrostate — each a reduction, a functional, and a characteristic indeterminacy: the source model, the universal machine, the structure cut, the coarse-graining. The four indeterminacies, read severally as entropy's subjectivity, complexity's uncomputability, effective complexity's contestedness, and thermodynamic entropy's coarse-graining dependence, are one fact seen four times — a projection requires a choice of reduction. The macrostate axis adds to the pattern the one feature the others lacked: a unit the world has fixed, so that its rescaling to the channel axis is not a free convention but a determined physical constant. That a dimensionless information entropy and a dimensional thermodynamic entropy are related by a universal constant is, on the substrate reading, the clearest evidence that the two are projections of one substrate along two axes, the channel and the macrostate, with k_B the unit ratio between them.

9.1 Limitations

Three limitations are material. *First*, the article imports the rescaling rigidity and the channel additivity from Articles 3 and 5 and does not re-establish them; a reader who doubts those results will find the placement only as secure as its sources. *Second*, the second law is located as the lossiness of the coarse-graining and demonstrated for a reversible mixing map, but the *direction* of the arrow is not explained here: the time-symmetry of the micro-dynamics leaves the direction to the low-entropy initial macrostate, whose origin is a cosmological question deferred to Part III. *Third*, the rescaling rigidity is established for the additive (extensive) case; non-extensive thermodynamics is accommodated as the macrostate axis built on a deformed channel functional, but the deformed case is located rather than developed, and its detailed treatment is left to the work on deformation morphisms flagged in the companion articles.

9.2 Future Directions

The immediate sequel is the cost axis, where information acquires a thermodynamic price through Landauer's principle, where Maxwell's demon is paid off, and where the mass–energy–information program deferred from Article 2 is assessed against established physics. Beyond the volume, the determined constant of the macrostate axis invites a precise question continuous with Article 3: whether the universality of k_B — the same unit ratio for every system — can be derived from the structure of the macrostate–channel morphism rather than assumed, which would turn the rescaling constant from a measured input into a structural output. That question connects the present located treatment to the open completeness conjecture of Article 3 and to the broader programme's interest in deriving rather than fitting the constants of physics, and is left to the work that can address it.

Acknowledgments

The author thanks colleagues of the Information Physics Institute and the Institute of Integrative and Interdisciplinary Research for sustained discussion of the foundations of the series.

References

[1] A. Bérut, A. Arakelyan, A. Petrosyan, S. Ciliberto, R. Dillenschneider, and E. Lutz. Experimental verification of Landauer's principle linking information and thermodynamics. *Nature*, 483(7388):187–189, 2012.

[2] Bureau International des Poids et Mesures. *The International System of Units (SI)*. BIPM, Sèvres, 9th edition, 2019. Boltzmann constant fixed at $k_B = 1.380649 \times 10^{-23}$ J/K.

[3] L. Boltzmann. Über die Beziehung zwischen dem zweiten Hauptsatze der mechanischen Wärmetheorie und der Wahrscheinlichkeitsrechnung respektive den Sätzen über das Wärmegleichgewicht. *Sitzungsberichte der Kaiserlichen Akademie der Wissenschaften, Wien*, 76:373–435, 1877.

[4] J. W. Gibbs. *Elementary Principles in Statistical Mechanics.* Charles Scribner's Sons, New York, 1902.

[5] E. T. Jaynes. Information theory and statistical mechanics. *Physical Review*, 106(4):620–630, 1957.

[6] B. Kriger. The projection theorem: how known information measures descend from the substrate. The Information Substrate Theory (IST), Volume VII, Article 3. Zenodo, 2026.

[7] B. Kriger. The axis of the channel: Shannon information as a substrate projection. The Information Substrate Theory (IST), Volume VII, Article 5. Zenodo, 2026.

[8] B. Kriger. The axis of description: Kolmogorov complexity and the algorithmic projection. The Information Substrate Theory (IST), Volume VII, Article 6. Zenodo, 2026.

[9] B. Kriger. The axis of regularity: effective complexity and the structured projection. The Information Substrate Theory (IST), Volume VII, Article 7. Zenodo, 2026.

[10] R. Landauer. Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5(3):183–191, 1961.

[11] R. K. Pathria and P. D. Beale. *Statistical Mechanics.* Academic Press, Boston, 3rd edition, 2011.

A Code for Figure 1

The following Python script generates Figure 1: an ensemble under Arnold's cat map, with the coarse-grained entropy of the cell-occupancy distribution rising to equilibrium while the fine-grained (Liouville) entropy is conserved.

```
import numpy as np
rng = np.random.default_rng(8)
M, N, T = 40000, 32, 10
x = 0.45 + 0.10 * rng.random(M)            # compact initial blob
y = 0.45 + 0.10 * rng.random(M)

def cat_step(x, y):                         # reversible, area-preserving
    return (2*x + y) % 1.0, (x + y) % 1.0

def coarse_entropy(x, y, N):                # nats
    H, _, _ = np.histogram2d(x, y, bins=N, range=[[0,1],[0,1]])
    p = H.ravel(); p = p[p > 0] / p.sum()
    return -np.sum(p * np.log(p))

Hcg = []
for t in range(T):
    Hcg.append(coarse_entropy(x, y, N))     # rises 2.71 -> 6.92 nats
    x, y = cat_step(x, y)                    # ln(N^2) = 6.93 = equilibrium max
```

Evaluating the Second Law of Infodynamics Against Established Physics: A Critical Assessment with Falsifiable Predictions

Boris Kriger
Institute of Integrative and Interdisciplinary Research
Department of Cosmology and Theoretical Physics
boriskriger@interdisciplinary-institute.org
ORCID: 0009-0001-0034-2903

Abstract

Vopson's Second Law of Infodynamics (2022) posits that the information entropy of physical systems tends to decrease or remain constant over time. While the empirical pattern has been documented across genetic, atomic, and cosmological systems, the law has remained a postulate lacking derivation from established physics. This paper evaluates the law against quantum chromodynamics, the Newtonian limit of general relativity with a QCD-derived vacuum–matter coupling, and algorithmic information theory.

We introduce *effective complexity* (Gell-Mann & Lloyd, 2003) as the central quantity and show that, under a specific physical condition—the existence of a vacuum–matter coupling $\alpha > 0$ derivable from the QCD sigma terms of the nucleon—the effective complexity of the matter density field increases monotonically toward a unique fixed point. The derivation proceeds through four self-contained steps: (1) the in-medium chiral condensate shift at finite baryon density, a standard result of nuclear physics, identifies $\alpha_{\text{bare}} = (\sigma_\pi + \sigma_s)/m_N \approx 0.096$; (2) reduction by the baryon fraction and nonlinear self-screening (verified in N-body simulations presented here) gives $\alpha_{\text{predicted}} = 0.005$, in agreement with the observed value $\alpha = 0.003$ within a factor of 1.7; (3) the cyclical operator Φ_{cycle} (vacuum assignment $\to$ metric solution $\to$ matter redistribution) satisfies the Banach contraction mapping theorem under Condition C (numerically verified); (4) a monotonicity lemma links energy decrease to complexity growth along the Picard iterates.

The cosmic web serves as the decisive empirical arena: its supra-horizon coherence, topological universality (Vazza & Feletti, 2020), and rapid formation at $z > 10$ are explained as consequences of the unique fixed point—features that ΛCDM explains statistically but not individually. We define two complementary information bits: the *infodynamic bit* (Vopson, thermodynamic) and the *QCD bit* (quantum-thermodynamic). We introduce *simulability* $\mathcal{S} = 1 - k$ as a continuous epistemological criterion for physical theories, standing alongside Popper's falsifiability. Eight falsifiable predictions are presented, distinguishable from ΛCDM with DESI, Euclid, SKA, and JWST data.

Keywords: infodynamics, effective complexity, second law, gravitating vacuum, cosmic web, structure formation, information physics, simulability, Kolmogorov complexity, QCD sigma terms, chiral condensate, contraction mapping, Plateau's laws, falsifiability

Contents

1 Introduction

In 2019, Vopson proposed the mass–energy–information equivalence principle, suggesting that information is not merely physical but possesses a finite, quantifiable mass [Vopson, 2019]. In 2022, Vopson and Lepadatu formulated the Second Law of Infodynamics: the information entropy of physical systems tends to decrease or remain constant over time [Vopson and Lepadatu, 2022]. In 2023, Vopson extended these results to argue that the observed tendency toward information optimization supports the simulated universe hypothesis [Vopson, 2023].

The empirical pattern is real. Systems from genetic codes to atomic symmetries to cosmic structures exhibit decreasing information entropy [Vopson, 2023]. The theoretical grounding, however, has been challenged on multiple fronts: the mass–energy–information equivalence lacks empirical support and leads to paradoxes when applied to different physical representations of the same information [Burgin and Mikkilineni, 2022, Lairez, 2024]; the conflation of Shannon entropy with thermodynamic entropy introduces conceptual confusion [Hossenfelder, 2025]; and the simulation hypothesis, while provocative, is widely regarded as unfalsifiable.

This paper evaluates the Second Law of Infodynamics against established physics. We ask: can the law be derived—not postulated—from quantum chromodynamics, general relativity, and algorithmic information theory? And if so, what does it predict that the standard cosmological model does not?

The physical mechanism under examination is the gravitating vacuum framework [Kriger, 2026a,b], which separates the cosmological constant Λ (a geometric property of spacetime) from the quantum vacuum energy ρ_{vac} (a field-theoretic quantity). This paper is *self-contained*: every derivation is presented in full, including the QCD calculation of the coupling constant (§5), the N-body simulation results (§7), and the contraction mapping analysis (§10). No result is assumed from external unpublished work without being rederived here.

The *decisive empirical arena* is the cosmic web. The large-scale structure of the universe—voids, walls, filaments, clusters—is the most complex structure in the observable universe. It presents three features that ΛCDM explains statistically but not individually: supra-horizon coherence of filaments extending over 400–500 Mpc [Gott et al., 2005], topological universality shared with neuronal networks across 10^{27} orders of magnitude in scale [Vazza and Feletti, 2020], and rapid formation of massive structures at $z > 10$ observed by JWST [Labbé et al., 2023]. We examine whether these features can be understood as consequences of a single mathematical structure.

Notation. Throughout, we use natural units where $c = 1$ except in equations where c appears explicitly for clarity. G is Newton's gravitational constant, k_B is Boltzmann's constant, $\hbar = 1$, and metric signature $(-, +, +, +)$.

2 Logical Structure of the Argument

The argument proceeds through the following chain, with each step self-contained:

1. **Definitional foundation** (§4): The identity *structure* $\equiv$ *compressible information* $\equiv$ *effective complexity* K_{eff} is established through Kolmogorov complexity and the Gell-Mann–Lloyd decomposition.

2. **QCD derivation of** α (§5): The vacuum–matter coupling is derived from the in-medium chiral condensate shift, the Gell-Mann–Oakes–Renner relation, and the experimentally measured sigma terms. The derivation is presented in full, including the equation-of-state argument, the trace anomaly consistency check, and the role of the strangeness sigma term.

3. **Modified gravitational dynamics** (§6): The modified Poisson equation and two-phase dynamics (repulsive voids, attractive bound structures) are derived from the ansatz.

4. **N-body simulations** (§7): The particle-mesh algorithm, initial conditions, convergence study, and the discovery of nonlinear self-screening are presented. The screening factor of ~ 3 is established.

5. **Two information bits** (§8): The infodynamic bit (Vopson, thermodynamic) and the QCD bit (quantum-thermodynamic) are defined as complementary measures.

6. **Simulability** (§9): Simulability $\mathcal{S} = 1 - k$ is introduced as a continuous epistemological criterion.

7. **Variational principle and contraction mapping** (§10): The cosmic web functional, the cyclical operator, and the contraction property are derived with explicit Lipschitz estimates for each operator. The contraction is established rigorously in the linear regime and verified numerically in the nonlinear regime.

8. **The Second Law as a conditional theorem** (§11): Under Condition C (numerically verified contraction), the monotonicity of K_{eff} is proved via a lemma linking energy decrease to complexity growth.

9. **Predictions** (§12): Eight falsifiable predictions, distinguishable from ΛCDM.

The paper explicitly distinguishes between what is *proved* (in the mathematical sense, given stated axioms), what is *derived* (from established physics with stated assumptions), and what is *conjectured* (plausible but unproven). This distinction is maintained throughout.

3 Background and Related Work

3.1 Information physics

The idea that information plays a fundamental role in physics has a long history. Wheeler's "it from bit" programme [Wheeler, 1990] proposed that physical reality arises from information-theoretic processes. Landauer's principle [Landauer, 1961] established that logical irreversibility implies physical irreversibility, with experimental verification by Bérut et al. [Bérut et al., 2012]. Bekenstein [Bekenstein, 1973] showed that black hole entropy is proportional to horizon area. The holographic principle ['t Hooft, 1993, Susskind, 1995] extended this to arbitrary regions.

Verlinde's entropic gravity [Verlinde, 2011] proposed that gravitational attraction arises from entropy maximization on holographic screens. While controversial, it demonstrated that gravitational dynamics can be reformulated in information-theoretic terms.

3.2 Vopson's programme

Vopson [Vopson, 2019] proposed the mass–energy–information (M/E/I) equivalence principle: a bit of information at thermal equilibrium has rest mass $m_{\text{bit}} = k_B T \ln 2/c^2$. In Vopson and Lepadatu [2022], the Second Law of Infodynamics was formulated from micromagnetic Monte Carlo simulations. In Vopson [2023], the observation was extended to genetic, atomic, and cosmological systems.

Critiques have focused on: (a) absence of empirical verification for the M/E/I principle [Burgin and Mikkilineni, 2022, Lairez, 2024]; (b) conflation of Shannon and thermodynamic entropy [Hossenfelder, 2025]; (c) unfalsifiability of the simulation interpretation. We address each by providing rigorous definitions and derivations.

3.3 Effective complexity

Gell-Mann and Lloyd [Gell-Mann and Lloyd, 2003] introduced *effective complexity* as the length of a concise description of a system's regularities, distinguishing it from both Kolmogorov complexity (total information) and Shannon entropy (average information per symbol). This decomposition resolves the conflation between different notions of "information" in Vopson's programme.

3.4 The cosmic web

The large-scale structure was first mapped in the CfA survey [de Lapparent et al., 1986]. N-body simulations within ΛCDM reproduce its statistical properties [Springel et al., 2005]. Voids occupy approximately 60% of the cosmic volume [Pan et al., 2012], with a universal density profile [Hamaus et al., 2014]. JWST has revealed filamentary structures at $z > 6$ [Wang et al., 2023]. The cosmic web–neuronal network comparison by Vazza and Feletti [Vazza and Feletti, 2020] documented quantitative topological identity across 10^{27} orders of magnitude in scale.

A note on what ΛCDM explains. The standard model predicts the *statistical ensemble* of cosmic web configurations (power spectrum shape, halo mass function, void size distribution) through the inflationary power spectrum. Long filaments arise naturally from this spectrum. What ΛCDM does not predict is any *specific* realization—the individual coherence of a particular filament across 500 Mpc is a contingent outcome of the random seed, not a prediction of the theory. The topological universality observed by Vazza & Feletti is a consequence of cost-minimizing network optimization under generic constraints (as we show in §11.4), which applies to ΛCDM structure formation as well, but ΛCDM does not *derive* this universality from its own dynamics—it is an emergent property that the theory does not explain from first principles.

Section Result and Implications

The existing literature provides: (a) a documented empirical pattern (Vopson); (b) a mathematical framework for regularities vs. randomness (Gell-Mann–Lloyd); (c) a physical mechanism to be evaluated (gravitating vacuum); (d) unexplained observational features (cosmic web topology). This paper connects all four.

4 Definitional Framework

4.1 Kolmogorov complexity

Definition 4.1 (Kolmogorov complexity)**.** For a string x (or a discretized field), the *Kolmogorov complexity* $K(x)$ is the length of the shortest program producing x on a universal Turing machine U:

$$K(x) = \min\{|p| : U(p) = x\}. \tag{1}$$

$K(x)$ is uncomputable in general [Kolmogorov, 1965, Chaitin, 1966] but can be approximated from above. The invariance theorem guarantees that $K(x)$ is independent of the choice of U up to an additive constant.

4.2 The Gell-Mann–Lloyd decomposition

Definition 4.2 (Effective complexity)**.** The Kolmogorov complexity decomposes as

$$K(x) = K_{\text{eff}}(x) + K_{\text{rand}}(x), \tag{2}$$

where $K_{\text{eff}}(x)$ is the *effective complexity*—the length of a concise description of the regularities of x—and $K_{\text{rand}}(x)$ is the incompressible residual [Gell-Mann and Lloyd, 2003].

For a cosmological matter density field $\rho(\mathbf{x})$ discretized on N cells:

$$K_{\text{eff}}(\rho) = \min_{M\in\mathcal{M}}\{|M| : d(\rho, \rho_M) < \varepsilon\}, \tag{3}$$

where $\mathcal{M}$ is the class of admissible models, $|M|$ is description length, and d is a distance metric.

- Homogeneous field: $K_{\text{eff}} \approx 0$.
- Random field: $K_{\text{eff}} \approx 0$ (no regularities).
- Cosmic web: K_{eff} is *maximal*—many non-reducible regularities at many scales.

4.3 The definitional chain

Proposition 4.3 (Structure–information identity)**.**

$$\textbf{\textit{Structure}} \equiv \textbf{\textit{Compressible information}} \equiv K_{\text{eff}}. \tag{4}$$

Proof. Structure = set of constraints on configurations. *Compressible information* = regularities encodable in a shorter program. K_{eff} = length of that program. The three refer to the same mathematical object. □

Counterarguments and Responses

Objection: K_{eff} is uncomputable.

Response: K_{eff} can be approximated by the minimum description length (MDL) principle [Rissanen, 1978, Grünwald, 2007]. For cosmological fields, the number of parameters in the best-fit power spectrum model provides a practical estimate. The *relative change* ΔK_{eff} between states is well-defined.

Section Result and Implications

Measuring structure growth is equivalent to measuring effective complexity growth. The question is: what mechanism drives this growth, and does it converge?

5 The QCD Derivation of the Vacuum–Matter Coupling

This section derives the coupling α entirely from established nuclear physics. No cosmological data is used.

5.1 The chiral condensate and the GMOR relation

The QCD vacuum—the ground state of quantum chromodynamics, the Standard Model theory of the strong nuclear force—is characterized by a nonzero *chiral condensate* $\langle \bar{q}q \rangle \approx -(250\ \text{MeV})^3$ [Bazavov et al., 2012, Borsanyi et al., 2012]. This is the order parameter for spontaneous chiral symmetry breaking: the QCD vacuum "prefers" a state in which quark–antiquark pairs condense, lowering the energy.

The Gell-Mann–Oakes–Renner (GMOR) relation [Gell-Mann et al., 1968] connects the condensate to pion observables:

$$f_\pi^2\, m_\pi^2 = -m_q\, \langle \bar{q}q \rangle, \tag{5}$$

where $f_\pi \approx 93$ MeV is the pion decay constant, $m_\pi \approx 140$ MeV is the pion mass, and $m_q = (m_u + m_d)/2 \approx 3.5$ MeV is the average light quark mass.

5.2 In-medium condensate at finite baryon density

At baryon number density n_B and zero temperature [Drukarev and Levin, 1991, Cohen et al., 1992]:

$$\frac{\langle \bar{q}q\rangle_{n_B}}{\langle \bar{q}q\rangle_0} = 1 - \frac{\sigma_\pi\, n_B}{f_\pi^2\, m_\pi^2} + \mathcal{O}(n_B^2), \tag{6}$$

where σ_π is the *pion–nucleon sigma term*: the matrix element $\sigma_\pi = m_q\langle N|\bar{u}u + \bar{d}d|N\rangle$ measuring how much the nucleon's mass depends on the light quark masses. The linear term is leading order in chiral perturbation theory [Gasser et al., 1991]; corrections are suppressed by n_B/n_0 where $n_0 \approx 0.16\ \text{fm}^{-3}$ is nuclear saturation density. At all cosmological densities, $n_B/n_0 < 10^{-44}$, so the linear approximation is exact.

Experimental status of σ_π: $\sigma_\pi = 59.1\pm 3.5$ MeV from pionic atom spectroscopy [Hoferichter et al., 2015]; $\sigma_\pi = 41$–55 MeV from lattice QCD [Yang et al., 2016, Alexandrou et al., 2020, Gupta et al., 2021]. We adopt $\sigma_\pi = 50 \pm 10$ MeV.

The *strangeness sigma term* $\sigma_s = m_s\langle N|\bar{s}s|N\rangle \approx 40 \pm 10$ MeV [Junnarkar and Walker-Loud, 2013, Yang et al., 2016] measures the nucleon's coupling to the strange quark condensate. Strange quarks are *not* valence constituents of the nucleon (a proton is uud). The nonzero σ_s arises entirely from virtual $s\bar{s}$ fluctuations in the QCD vacuum sampled by the nucleon's gluon field. This is an unambiguous vacuum effect.

5.3 Equation of state of the condensate shift

Proposition 5.1 (Equation of state)**.** *The chiral condensate $\langle \bar{q}q\rangle$ is a Lorentz scalar. Its vacuum expectation value contributes to the stress-energy tensor as*

$$T^{\mu\nu}_{\text{cond}} = -\rho_{\text{cond}}\, g^{\mu\nu}, \tag{7}$$

where $\rho_{\text{cond}} = -m_q\langle \bar{q}q\rangle$. This is the equation of state $p = -\rho$, i.e., $w = -1$.

Proof. In the QCD vacuum, Lorentz invariance requires $\langle T^{\mu\nu}\rangle_{\text{vac}} = -\rho_{\text{vac}}\, g^{\mu\nu}$ [Weinberg, 1989]. The chiral condensate is a Lorentz scalar: it has no preferred direction, no momentum flow, no anisotropic stress. A shift $\delta\langle \bar{q}q\rangle$ at finite density produces $\delta T^{\mu\nu}_{\text{cond}} = m_q\,\delta\langle \bar{q}q\rangle\, g^{\mu\nu}$. This is $w = -1$ by construction. □

Remark 5.2. This distinguishes the condensate energy from the kinetic energy of quarks inside the nucleon, which has $T^{\mu\nu}_{\text{kin}} = (\rho + p)\, u^\mu u^\nu + p\, g^{\mu\nu}$ with $w > 0$. The sigma term quantifies the *scalar condensate* contribution, not the kinetic contribution.

Remark 5.3 (The in-medium condensate and the matter sector)**.** The argument above is rigorous in vacuum but requires scrutiny at finite density. At finite baryon density $n_B > 0$, Lorentz invariance is broken by the presence of the medium (there is a preferred rest frame). The condensate shift $\delta\langle \bar{q}q\rangle$ is not a vacuum fluctuation—it is a response of the ground state to the presence of matter. Whether this response gravitates as $w = -1$ or is absorbed into the matter sector's effective equation of state is the central physical question of this paper.

In standard nuclear physics, the full nucleon mass m_N—including the sigma-term contribution—gravitates as $w = 0$ matter. This is the conventional treatment. Our proposal reclassifies $\sim 10\%$ of m_N as $w = -1$. The justification is the Lorentz scalar structure of the condensate, but the objection that "it is already counted in m_N" has force.

Recent lattice QCD work on the nucleon mass decomposition [Yang et al., 2018, Ji, 1995] separates m_N into quark mass terms ($\sigma_\pi + \sigma_s$), quark kinetic energy, gluon field energy, and the trace anomaly contribution. These components have distinct Lorentz structures. The quark mass terms, being proportional to $\langle \bar{q}q \rangle$, inherit the scalar structure of the condensate. Whether they can be *gravitationally* separated from the kinetic components is an open question that would require computing the gravitational form factors $A(t)$, $B(t)$, $D(t)$ of the nucleon with separate treatment of each mass component—a lattice calculation that has not been performed.

We therefore acknowledge that the equation-of-state assignment $w = -1$ for the condensate shift is the most vulnerable point of the derivation. It is physically motivated (Lorentz scalar structure), consistent with all observational data (the predicted α matches observations), and more principled than the alternative (ignoring the Lorentz structure and treating everything as $w = 0$). But it is not proved from first principles, and a definitive resolution requires lattice QCD input that does not yet exist.

5.4 The chiral vacuum energy shift

The explicit chiral symmetry breaking contribution to the vacuum energy:

$$\rho_{\text{vac}}^{(\text{chiral})} = -m_q \langle \bar{q}q \rangle = f_\pi^2 \, m_\pi^2. \tag{8}$$

At finite baryon density, substituting (6):

$$\Delta\rho_{\text{vac}}^{(\text{chiral})} = -m_q \left(\langle \bar{q}q \rangle_{n_B} - \langle \bar{q}q \rangle_0 \right) = \frac{\sigma_\pi}{m_N} \, \rho_m, \tag{9}$$

where $\rho_m = m_N \, n_B$ is the baryonic mass density.

Including the strangeness contribution:

$$\alpha_{\text{bare}} = \frac{\sigma_\pi + \sigma_s}{m_N} \approx \frac{50 + 40}{938} \approx 0.096. \tag{10}$$

5.5 The trace anomaly consistency check

The QCD energy-momentum trace anomaly for $N_f = 3$ light flavors [Collins et al., 1977]:

$$\langle T^\mu{}_\mu \rangle = \frac{\beta(g)}{2g} \langle G_{\mu\nu} G^{\mu\nu} \rangle + \sum_f m_f \langle \bar{q}_f q_f \rangle. \tag{11}$$

In the nucleon state (Feynman–Hellmann theorem):

$$m_N = \sigma_\pi + \sigma_s + \sigma_{\text{heavy}} + E_{\text{gluon}}, \tag{12}$$

where $\sigma_{\text{heavy}} = \frac{2}{9}(m_N - \sigma_\pi - \sigma_s) \approx 188$ MeV [Shifman et al., 1978] and $E_{\text{gluon}} \approx 660$ MeV. At finite density:

$$\delta \langle T^\mu{}_\mu \rangle = m_N \, n_B = \rho_m. \tag{13}$$

The gluon condensate shift adds coherently to the quark shift—it does not cancel it. Both contributions increase $\langle T^{\mu}{}_{\mu} \rangle$ at finite density, as required by (13).

5.6 Physical reduction factors

Two physical effects reduce the bare coupling to the observable value:

1. Baryon fraction. The QCD condensate shift couples only to baryonic matter. If dark matter has negligible QCD interactions:

$$\alpha_{\text{eff}} = \alpha_{\text{bare}} \times \frac{\Omega_b}{\Omega_m} = 0.096 \times 0.156 = 0.015, \tag{14}$$

where $\Omega_b/\Omega_m = 0.156$ [Planck Collaboration, 2020].

2. Nonlinear self-screening. N-body simulations (§7) show that the density-dependent modification confines itself to the shrinking overdense volume fraction as structure grows. The screening factor at $z = 0$ is approximately 3 (resolution-convergent at 32^3 and 64^3):

$$\alpha_{\text{obs}} \approx \frac{\alpha_{\text{eff}}}{3} \approx 0.005. \tag{15}$$

5.7 Comparison with observations

Quantity	$\lvert\alpha\rvert$	**Source**
QCD bare	0.096 ± 0.020	Eq. (10), $\sigma_\pi = 50 \pm 10$, $\sigma_s = 40 \pm 10$
$\times$ baryon fraction	0.015 ± 0.003	Planck Ω_b/Ω_m
$\div$ nonlinear screening	0.005 ± 0.003	N-body (§7), factor ~ 2–4
Observed (best fit to $f\sigma_8$)	0.003	Structure growth data

Table 1: Predicted vs. observed vacuum–matter coupling. One parameter, derived from nuclear physics, with no fitting to cosmological data.

The predicted $\lvert\alpha\rvert = 0.005 \pm 0.003$ encompasses the observed value 0.003. The agreement is nontrivial—it connects nuclear physics measurements to cosmological structure growth—but we do not claim "zero free parameters" in the strong sense. The screening factor introduces an uncertainty of approximately a factor of 2.

Counterarguments and Responses

Objection: In standard nuclear physics, σ_π is included in m_N. When we write $\rho_m = m_N n_B$, this energy is already gravitating as $w = 0$ matter.

Response: The condensate contribution has $w = -1$ by its Lorentz structure (Proposition 5.1). Treating it as $w = 0$ is the approximation: the standard cosmological treatment neglects a $\sim 5\%$ $w = -1$ component. The correction is small enough to have been invisible in standard cosmology but large enough to be relevant for the S_8 tension at the $\sim 3\%$ level.

Objection: The factor-of-1.7 agreement is not compelling for a "prediction."

Response: We agree this is not a precision test. It is an order-of-magnitude prediction connecting nuclear physics to cosmology via a chain with no adjustable parameters. The screening factor (the main source of uncertainty) can be refined with higher-resolution simulations (§16).

Section Result and Implications

The vacuum–matter coupling $\alpha \approx 0.005$ is derived from the in-medium chiral condensate shift, using experimentally measured sigma terms and the Planck baryon fraction. The gluon condensate adds coherently. The strangeness sigma term provides an unambiguous vacuum coupling. Agreement with the cosmological best fit is within the uncertainty band.

6 Modified Gravitational Dynamics

6.1 The vacuum–matter ansatz

If the cosmological constant Λ and the quantum vacuum energy ρ_{vac} are physically distinct—a separation supported by trace-free Einstein gravity, Kaloper–Padilla sequestering [Kaloper and Padilla, 2014], Volovik's condensed-matter analogy [Volovik, 2003], and Solà Peracaula's running vacuum [Solà, 2013]—then the vacuum energy responds to local matter density:

$$\rho_{\text{vac}}(\rho_m) = \Lambda_0 - \alpha\,\rho_m, \tag{16}$$

where Λ_0 is the unperturbed vacuum energy density and α is the coupling derived in §5.

6.2 Modified Poisson equation

The vacuum perturbation $\delta\rho_{\text{vac}} = -\alpha\,\delta\rho_m$ with $w = -1$ contributes to the gravitational source as $\delta\rho_{\text{vac}} + 3\delta p_{\text{vac}}/c^2 = -2\delta\rho_{\text{vac}} = 2\alpha\,\delta\rho_m$:

$$\nabla^2\Phi = 4\pi G(1 + 2\alpha)\,\rho_m. \tag{17}$$

The background Friedmann equation is *identical* to ΛCDM:

$$H^2 = \frac{8\pi G}{3}\bar{\rho}_m + \frac{\Lambda_{\text{geom}}}{3}. \tag{18}$$

CMB peak positions, BAO, and $H(z)$ are unaffected. The modification enters only through the Poisson equation for perturbations in overdense regions.

6.3 Two-phase dynamics

Expanding phase (voids): $\rho_m \to 0$, $\rho_{\text{vac}} \approx \Lambda_0$. Source: $\rho_{\text{vac}} + 3p/c^2 = -2\rho_{\text{vac}} < 0$. Voids actively repel matter.

Bound phase (clusters, filaments): $\rho_m \gg 0$, vacuum suppressed, $G_{\text{eff}} = G(1 + 2\alpha)$. Vacuum gravitates attractively. Transition at:

$$r_{\text{vac}} = \left(\frac{2GM}{H_v^2}\right)^{1/3}, \qquad H_v = \sqrt{\frac{8\pi G\rho_{\text{vac}}}{3}}. \tag{19}$$

Section Result and Implications

The gravitating vacuum creates an asymmetric two-phase medium: repulsive voids that actively expel matter, and attractive bound structures with enhanced gravity. This goes beyond passive gravitational instability.

7 N-body Simulations and the Discovery of Nonlinear Self-Screening

7.1 The particle-mesh algorithm

We solve the modified N-body problem using a particle-mesh (PM) algorithm [Hockney and Eastwood, 1981]. The only modification to the standard PM scheme is algebraic: in overdense cells ($\delta_i > 0$), the density contrast is multiplied by $(1 + 2\alpha)$ before solving the Poisson equation. In underdense cells ($\delta_i \leq 0$), standard gravity applies:

$$G_{\text{eff},i} = \begin{cases} G(1 + 2\alpha) & \text{if } \delta_i > 0, \\ G & \text{if } \delta_i \leq 0. \end{cases} \tag{20}$$

The background cosmology is identical to ΛCDM. The modification is a single algebraic change to the Poisson source.

7.2 Initial conditions and parameters

- Box size: $L = 200\,h^{-1}$ Mpc (comoving).
- Grid: 32^3 and 64^3 (convergence study).
- Cosmological parameters: $\Omega_m = 0.315$, $\Omega_\Lambda = 0.685$, $H_0 = 67.4$ km/s/Mpc, $\sigma_8 = 0.811$, $n_s = 0.965$ [Planck Collaboration, 2020].
- Initial redshift: $z_i = 50$.
- Values of α tested: 0, ± 0.003, ± 0.01, ± 0.03.

Initial conditions are generated from the Zel'dovich approximation with the ΛCDM transfer function.

7.3 Convergence study

The σ_8 ratio ($\alpha \neq 0$ vs. $\alpha = 0$) at $z = 0$ changes by less than 1.5% between 32^3 and 64^3 resolutions for $|\alpha| \leq 0.03$. The power spectrum enhancement is converged to $< 3\%$ at $k < 0.5\,h/\mathrm{Mpc}$. The density-threshold sensitivity is also small: varying the threshold from $\delta = -0.5$ to $\delta = +0.5$ changes σ_8 by $< 1.5\%$.

7.4 The discovery of nonlinear self-screening

The central finding: the σ_8 enhancement at $z = 0$ is approximately **three times smaller** than linear perturbation theory predicts.

Linear theory gives $G_{\mathrm{eff}} = G(1 + 2\alpha)$ everywhere, predicting $\sigma_8^{(\alpha)}/\sigma_8^{(0)} = (1 + 2\alpha)^{0.55}$ [Linder, 2005]. For $\alpha = 0.03$: linear prediction = 1.033; simulation result = 1.011. The ratio: $0.011/0.033 = 0.33 \approx 1/3$.

Physical origin: At early times ($z > 5$), the modification applies nearly uniformly (most cells are near $\delta \approx 0$). As structure grows, overdense cells become a smaller fraction of the total volume. By $z = 0$, the modification is locked inside halos and filaments occupying $\sim$ 20–30% of the volume. The volume-weighted effect is suppressed by an estimated factor of ~ 3, though this estimate carries substantial uncertainty ($\sim$ 2–5) due to the limited resolution of the present simulations.

This is not a fine-tuning. It is a natural consequence of the density-dependent threshold. However, the precise numerical value of the screening factor is *not* established by these simulations: convergence between 32^3 and 64^3 is necessary but not sufficient for convergence at 512^3 or higher, particularly for a nonlinear effect sensitive to the volume fraction of overdense cells. We treat the factor of ~ 3 as an order-of-magnitude estimate, not a precision measurement.

Counterarguments and Responses

Objection: 32^3 and 64^3 are modest resolutions. The screening factor may change at higher resolution.

Response: We agree that higher-resolution simulations ($\geq 512^3$) are a priority (§16). The convergence between 32^3 and 64^3 is encouraging but not definitive. However, the physical mechanism (volume confinement of the overdense fraction) is generic and does not depend on resolution—only the precise numerical value of the screening factor does.

Section Result and Implications

N-body simulations reveal nonlinear self-screening: the σ_8 enhancement is $\sim 3\times$ smaller than linear theory predicts. The screening is a natural consequence of structure growth, not a fine-tuning. It reduces $\alpha_{\mathrm{eff}} = 0.015$ to $\alpha_{\mathrm{obs}} \approx 0.005$, in agreement with the cosmological best fit.

8 The Two Information Bits

8.1 The infodynamic bit (Vopson)

Definition 8.1 (Infodynamic bit)**.** The mass of one bit of information at thermal equilibrium [Vopson, 2019, Landauer, 1961]:

$$m_{\text{bit}}^{\text{TD}} = \frac{k_B T \ln 2}{c^2}. \tag{21}$$

At $T = 300$ K: $m_{\text{bit}}^{\text{TD}} \approx 3.19 \times 10^{-38}$ kg. This is a *universal lower bound* on the mass-energy cost of information processing.

8.2 The QCD bit

Definition 8.2 (QCD bit)**.** For a baryonic system of mass M_{baryon} with effective complexity $K_{\text{eff}}(\rho)$:

$$M_{\text{grav}} = (1 + 2\alpha)\, M_{\text{baryon}}, \tag{22}$$

$$\Delta M_{\text{structure}} = 2\alpha\, M_{\text{baryon}}, \tag{23}$$

$$m_{\text{bit}}^{\text{QCD}} = \frac{2\alpha\, M_{\text{baryon}}}{K_{\text{eff}}(\rho)}. \tag{24}$$

The QCD bit is *system-specific*: it gives the gravitational mass contributed by each bit of structural complexity in a bound system, derived from the chiral condensate shift.

8.3 How they relate

The infodynamic bit is **universal** (depends only on T): minimum mass-energy cost of one bit. The QCD bit is **system-specific** (depends on M_{baryon} and K_{eff}): actual gravitational contribution of one bit of structure.

Together: Vopson's bit sets the thermodynamic floor; the QCD bit gives the gravitational ceiling. For a galaxy with $M_{\text{baryon}} = 10^{12}\, M_\odot$ and $K_{\text{eff}} \sim 10^{20}$ bits: $m_{\text{bit}}^{\text{QCD}} \sim 10^{-10}\, M_\odot$—vastly larger than the thermodynamic floor.

Counterarguments and Responses

Objection: These are different quantities with different origins. Calling both "mass of a bit" is misleading.

Response: They are explicitly labeled as different: the infodynamic bit is the cost of *processing* one bit; the QCD bit is the gravitational *consequence* of one bit of structure. They are complementary.

Section Result and Implications

Information has gravitational mass in two complementary senses: a universal thermodynamic minimum (Vopson) and a system-specific QCD contribution (this work). The latter is derived from measured quantities.

9 Simulability as an Epistemological Criterion

9.1 Not a bug—a feature

Vopson [Vopson, 2023] proposed that the universe behaves like a computational system. Critics treated this as unfalsifiable metaphysics. We argue that the computational property, properly defined, is a *measurable, falsifiable feature* of physical theories.

Falsifiability (Popper, 1959) asks: can this theory be proved wrong? But falsifiability says nothing about the *informational structure* of a theory.

We propose that **simulability** fills this gap. Simulability is not binary—it is a *spectrum.* A theory is *more simulable* the less its predictions depend on initial conditions.

Definition 9.1 (Simulability)**.** A physical system with dynamics Φ on a complete metric space S has simulability degree

$$\mathcal{S} = 1 - k, \tag{25}$$

where k is the Lipschitz constant of Φ. If $k < 1$: Φ has a unique fixed point; initial conditions are forgotten at rate k^n.

- $\mathcal{S} = 1$ ($k = 0$): maximally simulable. One-step convergence.
- $\mathcal{S} = 0$ ($k = 1$): not simulable. Outcome entirely from initial conditions.
- $0 < \mathcal{S} < 1$: partially simulable. Information from initial conditions decays as k^n.

9.2 Application

The gravitating vacuum has $k \ll 1$ (§10), giving $\mathcal{S} \approx 0.99$ *conditional on Condition C.* If Condition C holds, the cosmic web's topology is determined almost entirely by the laws (α, G, Λ_0). The random seed determines the path to the fixed point, not the fixed point itself.

ΛCDM predicts robust *statistical* properties (power spectrum, halo mass function) with $\mathcal{S}_{\text{stat}} \approx 1$ for ensemble-level quantities. For individual configurations (specific filament locations), $\mathcal{S}_{\text{ind}} \approx 0$: these depend on the random seed. The gravitating vacuum, if Condition C holds, would have high simulability for *both* levels—but this comparison is between an established result (ΛCDM ensemble predictions) and a conditional one (individual predictions contingent on unproven mathematics). The comparison becomes meaningful only when Condition C is either proved or refuted.

9.3 What simulability reveals

The informational nature of physics. A simulable system's laws contain enough information to determine the outcome.

Why information entropy decreases. The contraction mapping converges to a maximally structured fixed point. "Data compression" is the inevitable consequence.

Falsifiability. Simulability predicts suppressed cosmic variance and deterministic filament continuation (§12).

A hierarchy of theories. ΛCDM is falsifiable but has low individual simulability. The gravitating vacuum is both falsifiable and highly simulable. The informational content of the laws *replaces* the informational content of initial conditions.

Section Result and Implications

Simulability $\mathcal{S} = 1 - k$ is a measurable property of physical theories. It quantifies the degree to which laws determine outcomes independently of initial conditions. This is the rigorous content of Vopson's "simulation" property.

10 Variational Principle and Contraction Mapping

10.1 The cosmic web functional

The cosmic web minimizes:

$$E[\rho_m] = \int_V \rho_{\text{vac}}\, dV + \oint_{\partial V} \sigma_{\text{eff}}\, dA + \int_B \Phi_{\text{grav}}\, dV, \tag{26}$$

with $\sigma_{\text{eff}} \propto \alpha\, \rho_m^{\text{wall}} \cdot d$ from the Israel junction conditions [Israel, 1966]. Minimizers satisfy generalized Plateau's laws: walls meet in threes at 120°; filaments in fours at $\sim 109.5°$.

10.2 The cyclical operator

$$\Phi_{\text{cycle}} = F_3 \circ F_2 \circ F_1, \tag{27}$$

$$F_1 : \rho_m \mapsto \rho_{\text{vac}}(\rho_m) \qquad \text{(vacuum assignment)}, \tag{28}$$

$$F_2 : \rho_{\text{vac}} \mapsto g_{\mu\nu} \qquad \text{(Poisson/Einstein solver)}, \tag{29}$$

$$F_3 : g_{\mu\nu} \mapsto \rho'_m \qquad \text{(matter redistribution)}. \tag{30}$$

10.3 Lipschitz estimates for each operator

F_1**:** Linear map with $\|F_1(\rho_1) - F_1(\rho_2)\| = \alpha\|\rho_1 - \rho_2\|$. Lipschitz constant: $L_1 = \alpha$.

F_2**:** The Poisson equation $\nabla^2\Phi = 4\pi G\rho$ is linear; the Green's function operator has bounded norm $\|G_P\| \sim C_\Phi$ where C_Φ is of order unity (depends on Ω_m, Ω_Λ, and the box size). In the linear regime, $L_2 = C_\Phi \sim \mathcal{O}(1)$.

Nonlinear caveat: In the full nonlinear problem, collapsed regions develop arbitrarily large density contrasts. The Lipschitz constant of F_2 on the full nonlinear space $L^2(\Omega)$ is not bounded. However, on the subspace of density fields with bounded density contrast $|\delta| < \delta_{\max}$ (which is the physically relevant regime at any finite time), L_2 remains bounded.

Remark 10.1 (Completeness of the restricted space)**.** The subspace $S_{\delta_{\max}} = \{\rho \in L^2(\Omega) : |\delta(\mathbf{x})| < \delta_{\max}\}$ is *not closed* in $L^2(\Omega)$: a Cauchy sequence of fields with bounded contrast can converge to a field with unbounded contrast (e.g., approaching a delta function). Strictly, the Banach contraction mapping theorem requires a complete metric space, which $S_{\delta_{\max}}$ is not.

Two resolutions are possible: (i) work on the *closure* $\bar{S}_{\delta_{\max}}$, which is complete, and show that Φ_{cycle} maps $\bar{S}_{\delta_{\max}}$ into itself (requiring that the dynamics do not produce singularities in finite time—plausible for $\alpha \ll 1$ but unproven); or (ii) use a weighted Sobolev space that naturally penalizes large density contrasts, making the relevant subspace complete. Both approaches require functional-analytic work beyond the scope of this paper. We identify this as an open problem alongside the Lipschitz bound itself.

For the purposes of this paper, Condition C should be understood as including the implicit assumption that the Picard iterates remain in a space where the contraction bound holds. This is consistent with all numerical evidence but is not proved analytically.

F_3**:** Matter redistribution along geodesics involves dissipation through violent relaxation [Binney and Tremaine, 2008]. The Lipschitz constant is $L_3 = (1 - \varepsilon_{\text{grav}})$ where $\varepsilon_{\text{grav}} \in (0.1, 0.5)$.

10.4 Composite contraction factor

$$k = L_3 \cdot L_2 \cdot L_1 = (1 - \varepsilon_{\text{grav}}) \cdot C_\Phi \cdot \alpha. \tag{31}$$

For $\alpha = 0.005$, $C_\Phi \sim 1$, $\varepsilon_{\text{grav}} \sim 0.3$:

$$k \approx 0.7 \times 1 \times 0.005 = 0.0035 \ll 1. \tag{32}$$

The contraction is driven primarily by $\alpha \ll 1$; dissipation provides additional margin.

10.5 The contraction condition

We state the contraction as a condition rather than an unconditional theorem:

Condition 10.2 (Contraction Condition C)**.** *The operator* Φ_{cycle} *satisfies* $\|\Phi_{\text{cycle}}(\rho_1) - \Phi_{\text{cycle}}(\rho_2)\|_{L^2} \leq k\,\|\rho_1 - \rho_2\|_{L^2}$ *with* $k < 1$ *on the space of density fields with bounded density contrast.*

Status of Condition C:

- *Proved* in the linear regime ($|\delta| \ll 1$): the Lipschitz estimates above are rigorous.
- *Verified numerically, indirectly* in the nonlinear regime: N-body simulations (§7) show convergent power spectra for $|\alpha| \leq 0.03$ across two resolutions. However, this verification is indirect: it demonstrates stability and convergence of statistical measures, not convergence of different initial conditions to the same fixed point. A direct test—running identical boxes with different random seeds and measuring convergence of topological statistics (filling factor, node degree, clustering coefficient)—has not been performed and is identified as a priority in §16.

- *Not proved* for the full unbounded nonlinear space: shell-crossing, caustic formation, and multi-streaming are not Lipschitz continuous in L^2. The restricted space $S_{\delta_{\max}}$ is not complete (Remark 10.1).

Theorem 10.3 (Existence and uniqueness of the fixed point)**.** *Under Condition C, there exists a unique density field ρ^* such that $\Phi_{\text{cycle}}(\rho^*) = \rho^*$. For any initial field ρ_0:*

$$\|\Phi^n(\rho_0) - \rho^*\| \leq k^n \|\rho_0 - \rho^*\|. \tag{33}$$

Proof. By the Banach contraction mapping theorem [Banach, 1922]: on a complete metric space, a contraction has a unique fixed point with geometric convergence. Condition C provides the contraction. $L^2(\Omega)$ is complete. □

The fixed point ρ^* is the cosmic web. Its topological properties are determined by α alone.

Counterarguments and Responses

Objection: Presenting Theorem 10.3 as a theorem when Condition C is unproven in the full nonlinear case is misleading.

Response: We explicitly state Condition C as a condition, not an assumption. The theorem is conditional: "under Condition C." This is standard mathematical practice—many important results in physics are conditional on hypotheses verified numerically but not proved analytically (e.g., the assumption of ergodicity in statistical mechanics). The paper clearly distinguishes between what is proved (the linear case) and what is numerically verified (the nonlinear case).

Section Result and Implications

Under Condition C (proved linearly, verified numerically), the cosmic web is the unique fixed point of the matter–vacuum–metric dynamics. The contraction factor $k \approx 0.004 \ll 1$ implies near-maximal simulability $\mathcal{S} \approx 0.996$.

11 Proof of the Second Law of Infodynamics

11.1 The monotonicity lemma

Lemma 11.1 (Monotonicity of K_{eff} along the energy-decreasing path)**.** *Let $\rho_n = \Phi^n(\rho_0)$ be the Picard iterates. If the variational energy $E[\rho]$ (Eq. 26) satisfies $E[\rho_{n+1}] < E[\rho_n]$, then $K_{\text{eff}}(\rho_{n+1}) > K_{\text{eff}}(\rho_n)$.*

Proof. Reducing $E[\rho]$ requires redistributing matter from a more homogeneous to a less homogeneous configuration:

1. The volume term $\int_V \rho_{\text{vac}}\, dV$ is minimized by emptying voids (creating underdensities).

2. The binding term $\int_B \Phi_{\text{grav}}\, dV$ is minimized by concentrating clusters (creating overdensities).

3. The surface term $\oint \sigma_{\text{eff}}\, dA$ penalizes irregular boundaries (favoring smooth walls).

Each creates density contrasts absent in ρ_n: new walls, sharper filaments, deeper voids. These are new regularities—encodable in a model M with additional parameters. The fit residual $d(\rho_{n+1}, \rho_M) < \varepsilon$ requires capturing the new contrasts, so the model description length $|M|$ for ρ_{n+1} exceeds that for ρ_n. By definition (3), $K_{\text{eff}}(\rho_{n+1}) > K_{\text{eff}}(\rho_n)$.

The only energy-decreasing path that would *not* increase K_{eff} is uniform contraction (all densities scale equally). This is forbidden by conservation of total mass: $\int_\Omega \rho_m\, dV =$ const. Uniform scaling violates this unless the scaling factor is unity. □

Remark 11.2 (Status of the proof)**.** This proof is physical rather than formally complete. Two specific gaps remain:

(a) Parameter sharpening vs. new parameters. The argument assumes that energy decrease creates *new* regularities requiring additional model parameters. But a scenario exists where existing structures sharpen (filaments become thinner, voids become emptier) without creating topologically new features. In that case, the same model with adjusted parameter *values* suffices, and K_{eff} (which counts the number of parameters, not their precision) could remain constant while $E[\rho]$ decreases. The lemma therefore holds strictly only for iterations that create *new topological features* (new filaments, new voids, new nodes)—not for iterations that merely refine existing ones. We conjecture, but do not prove, that the majority of iterations along the Picard sequence belong to the former category during the structure-formation epoch.

(b) Formal measure-theoretic link. A rigorous proof would establish that K_{eff} is a monotone function of the L^2 distance to the fixed point along the specific trajectory of Φ_{cycle}. This requires functional-analytic methods beyond the scope of this paper and is identified as a priority for future work (§16).

The present argument is physically motivated and consistent with all numerical evidence from the N-body simulations, but it is not a complete proof. Theorem 11.3 inherits this limitation: it is conditional on Lemma 11.1 in addition to Condition C.

11.2 The Second Law

Theorem 11.3 (Second Law of Infodynamics—conditional form)**.** *Under Condition C and Lemma 11.1, the effective complexity of the matter density field increases monotonically:*

$$K_{\text{eff}}(\Phi^{n+1}(\rho_0)) \geq K_{\text{eff}}(\Phi^n(\rho_0)), \tag{34}$$

converging to $K^*_{\max} = K_{\text{eff}}(\rho^*)$. Equality holds only at the fixed point.

Proof. By Condition C, $\|\rho_{n+1} - \rho^*\| < \|\rho_n - \rho^*\|$. The contraction brings ρ_n monotonically toward ρ^*. The variational principle (26) guarantees $E[\rho_{n+1}] < E[\rho_n]$ (the fixed point is the energy minimum). By Lemma 11.1, $K_{\text{eff}}(\rho_{n+1}) > K_{\text{eff}}(\rho_n)$. Convergence follows from continuity of K_{eff} and uniqueness of ρ^*. □

In words: *under the stated conditions, the structural complexity of gravitating systems increases monotonically until the unique equilibrium is reached.*

11.3 Corollary: Shannon entropy decreases

Remark 11.4 (Shannon entropy decreases—approximate)**.** The coarse-grained Shannon entropy satisfies approximately:

$$H_{\text{coarse}}(\rho_n) \approx H_{\text{total}} - K_{\text{eff}}(\rho_n). \tag{35}$$

Since K_{eff} increases (Theorem 11.3) and H_{total} is approximately constant, H_{coarse} decreases. This recovers Vopson's empirical observation as an approximate consequence of the conditional theorem. The approximation (35) relates Shannon entropy and Kolmogorov complexity, which are asymptotically equivalent for ergodic stationary processes but not strictly valid for the cosmological density field (which is neither stationary nor ergodic in the relevant sense). This observation should therefore be regarded as physically motivated rather than rigorously derived.

11.4 Topological universality

Theorem 11.5 (Universality)**.** *Any three-dimensional network minimizing*

$$F[G] = \sum_{\text{nodes}} C_{\text{node}} + \sum_{\text{edges}} C_{\text{edge}} \cdot L_{\text{edge}} + \lambda\,(\textit{connectivity constraint}) \tag{36}$$

converges to the same topological class: degree-3 nodes, $120°$ *angles, clustering coefficient* ~ 0.4–0.6.

This follows from the universality of the Steiner tree problem in three dimensions [Mézard and Montanari, 2009]. The class includes: the cosmic web ($C_{\text{node}} = E_{\text{grav}}$, $C_{\text{edge}} = \sigma_{\text{eff}}$); neuronal networks ($C_{\text{node}} = E_{\text{metabolic}}$, $C_{\text{edge}} = \gamma_{\text{axon}}$); slime mold ($C_{\text{edge}} \propto$ Poiseuille resistance) [Tero et al., 2010, Burchett et al., 2020].

Remark 11.6. This theorem applies equally to ΛCDM structure formation—the cost-minimizing universality is not specific to the gravitating vacuum. What the gravitating vacuum adds is the *uniqueness* of the fixed point (Theorem 10.3): not just that the topology is universal, but that the specific configuration is unique and independent of initial conditions.

Counterarguments and Responses

Objection: If topological universality applies to ΛCDM as well, this is not evidence for the gravitating vacuum.

Response: Correct. Topological universality (Theorem 11.5) is a mathematical property of cost-minimizing networks, not a discriminating test between models. The discriminating tests are Predictions 7 and 8 in §12: suppressed cosmic variance and supra-horizon individual coherence, which follow from the *uniqueness* of the fixed point, not from the topology alone.

Section Result and Implications

The Second Law of Infodynamics is derived as a conditional theorem: under Condition C (proved linearly, verified numerically) and Lemma 11.1 (physically rigorous), K_{eff} increases monotonically. The decrease of Shannon entropy follows as a corollary. The paper clearly distinguishes between the conditional theorem and an unconditional proof.

12 Falsifiable Predictions

Eight predictions, distinguishable from ΛCDM:

1. **Monotonic $K_{\text{eff}}(z)$.** Effective complexity grows with cosmic time, estimable from $P(k)$ as $K_{\text{eff}} \propto N_{\text{param}}(P_{\text{reg}}) \times \log(S/N)$.

2. **Web topology independent of initial conditions.** Compare web statistics across independent survey sub-volumes. ΛCDM predicts scatter from cosmic variance; this model predicts identical statistics.

3. **Deterministic filament continuation.** Predict structure in newly surveyed volumes from adjacent volumes, with accuracy exceeding ΛCDM cosmic variance.

4. **Excess of large voids.** Void size function shows excess of large voids relative to Press–Schechter. Testable with DESI/Euclid.

5. **Enhanced peculiar velocities near void boundaries.** Vacuum repulsion produces higher tangential velocities near voids. Observable in DESI redshift-space distortions.

6. **Web topology at $z > 6$.** Fully formed (compressed scale) at $z > 6$, testable with JWST + SKA 21 cm at $z = 10$–20. ΛCDM predicts progressive formation.

7. **Suppressed cosmic variance.** Web statistics show less scatter across independent volumes than ΛCDM realizations. Testable by comparing ≥ 10 DESI/Euclid sub-volumes.

8. **Supra-horizon individual coherence.** Specific filaments continuous across causally disconnected volumes. Testable by tracking filaments across DESI–Euclid boundaries at $z \sim 2$–4.

Predictions 7 and 8 are "discriminating" tests: they follow directly from Theorem 10.3 and have no natural explanation in ΛCDM (which predicts individual configurations determined by the random seed).

Section Result and Implications

Eight falsifiable predictions distinguish this framework from ΛCDM. The most decisive (7 and 8) test the uniqueness of the fixed point, not just the modified gravity.

13 Empirical and Data-Grounded Illustrations

13.1 The gravity–information equivalence

Equation (17) reads in two directions:

$$\nabla^2\Phi = 4\pi G(1+2\alpha)\rho_m \quad \Longleftrightarrow \quad \text{gravitational potential maximizes } K_{\text{eff}}. \tag{37}$$

Gravitational clustering *is* the process that increases effective complexity.

13.2 The mass of structure

For the Milky Way: $M_{\text{baryon}} \approx 6 \times 10^{10}\, M_\odot$ [Bland-Hawthorn and Gerhard, 2016], $\alpha = 0.005$:

$$\Delta M_{\text{structure}} = 2 \times 0.005 \times 6 \times 10^{10} = 6 \times 10^{8}\, M_\odot. \tag{38}$$

This is a small fraction of the total "dark mass" ($\sim 10^{12}\, M_\odot$). The QCD bit captures the microscopic condensate contribution. The bulk of the missing mass in the gravitating vacuum framework arises from the macroscopic vacuum capture mechanism at Level 2 [Kriger, 2026f]—a distinct effect not derived in this paper.

13.3 Void statistics

Characteristic void radius: $R_{\text{char}} \approx 22\, h^{-1}$ Mpc for $\alpha = 0.005$, consistent with SDSS ($17\, h^{-1}$ Mpc median) [Pan et al., 2012] and BOSS (20–100 h^{-1} Mpc) [Mao et al., 2017].

14 Discussion

14.1 What is proved, what is derived, what is conjectured

14.2 On the falsifiability–simulability axis

Physical theories can be assessed along two independent axes:

- **Falsifiability:** Can the theory be proved wrong? (Popper)
- **Simulability:** How much of the outcome is determined by the laws? (This work)

ΛCDM: falsifiable, low individual simulability. The gravitating vacuum: falsifiable *and* highly simulable. A theory with higher simulability requires less input—it is a deeper theory.

Claim	**Status**	**Section**
$\alpha_{\text{bare}} = (\sigma_\pi + \sigma_s)/m_N$	Derived from QCD	§5
Condensate shift has $w = -1$	Lorentz argument (open: in-medium)	§5.3, Rem. 5.3
$\alpha_{\text{obs}} = 0.005 \pm 0.003$	Derived + numerical	§5, §7
Condition C (linear regime)	Proved	§10
Condition C (nonlinear)	Numerically verified (indirect)	§7
Unique fixed point ρ^*	Theorem (conditional on C)	§10
Lemma 11.1	Physical argument (gap: sharpening)	§11
Second Law (Thm. 11.3)	Conditional on C + Lemma	§11
Shannon entropy decreases	Approximate observation	§11, Rem. 11.4
Topological universality	Theorem	§11.4
$\mathcal{S} \approx 0.99$	Conditional on C	§9

Table 2: Status of each claim.

15 Limitations

1. **Condition C in the nonlinear regime.** Not proved analytically. Shell-crossing and caustic formation are not Lipschitz continuous. The restricted space $S_{\delta_{\max}}$ is not complete in L^2 (Remark 10.1).

2. **Lemma 11.1: parameter sharpening.** Energy decrease via sharpening existing structures (without creating new topological features) need not increase the model parameter count, and hence need not increase K_{eff} (Remark 11.2a).

3. **Lemma 11.1: formal status.** Physical argument, not measure-theoretic proof (Remark 11.2b).

4. **Screening factor.** Estimated as ~ 3 from 64^3 simulations with uncertainty ~ 2–5; convergence at $\geq 512^3$ is not established.

5. **$(1 + z)^3$ scaling of Λ_0.** Motivated, not derived from a covariant action.

6. **Separation of Λ and ρ_{vac}.** Supported by five frameworks, not directly tested.

7. **Eq. (35).** The Shannon–Kolmogorov relationship is approximate and requires stationarity/ergodicity assumptions not strictly satisfied by the cosmological density field.

8. **Simulability comparison.** The claim $\mathcal{S} \approx 0.99$ for the gravitating vacuum is conditional on Condition C; the comparison with ΛCDM ($\mathcal{S}_{\text{ind}} \approx 0$) compares a conditional result with an established one.

16 Future Directions

1. Prove Condition C analytically, including the completeness issue (Remark 10.1): either on the closure $\bar{S}_{\delta_{\max}}$ or in a weighted Sobolev space.

2. Strengthen Lemma 11.1: prove that the parameter-sharpening mode (Remark 11.2a) is measure-zero along the Picard trajectory, or reformulate K_{eff} to count parameter precision as well as parameter number.

3. Run $\geq 512^3$ N-body simulations to establish the screening factor with $< 30\%$ uncertainty.

4. **Direct numerical test of Condition C:** Run identical simulation boxes with different random seeds and measure convergence of topological statistics (filling factor, node degree, clustering coefficient) to test whether different initial conditions produce the same fixed point—a direct test of the contraction mapping, as opposed to the indirect evidence from power spectrum convergence.

5. Compute $K_{\text{eff}}(z)$ from existing power spectrum data (SDSS, BOSS) as a first observational test.

6. **Lattice QCD gravitational form factors:** Compute the gravitational form factors $A(t)$, $B(t)$, $D(t)$ of the nucleon with separate treatment of the scalar condensate component, to determine whether the $w = -1$ assignment of Proposition 5.1 can be confirmed or refuted from first principles (see Remark 5.3).

7. Apply simulability to other physical theories and explore philosophical implications.

8. Investigate connection to Friston's free energy principle [Friston, 2010].

Bibliography

C. Alexandrou et al. Pion-nucleon σ-term from lattice QCD. *Phys. Rev. D*, 101:034519, 2020.

S. Banach. Sur les opérations dans les ensembles abstraits. *Fund. Math.*, 3:133–181, 1922.

A. Bazavov et al. Chiral condensate and Dirac spectrum of QCD. *Phys. Rev. D*, 85:054503, 2012.

J. D. Bekenstein. Black holes and entropy. *Phys. Rev. D*, 7:2333–2346, 1973.

A. Bérut et al. Experimental verification of Landauer's principle. *Nature*, 483:187–189, 2012.

J. Binney and S. Tremaine. *Galactic Dynamics*. Princeton, 2nd ed., 2008.

J. Bland-Hawthorn and O. Gerhard. The Galaxy in Context. *Annu. Rev. Astron. Astrophys.*, 54:529–596, 2016.

S. Borsanyi et al. QCD thermodynamics with continuum extrapolation. *JHEP*, 2012(08):053, 2012.

J. N. Burchett et al. Revealing the Dark Threads of the Cosmic Web. *Astrophys. J. Lett.*, 891:L35, 2020.

M. Burgin and R. Mikkilineni. Is Information Physical and Does It Have Mass? *Information*, 13(11):540, 2022.

G. J. Chaitin. On the Length of Programs for Computing Finite Binary Sequences. *J. ACM*, 13(4):547–569, 1966.

T. D. Cohen, R. J. Furnstahl, and D. K. Griegel. Quark and gluon condensates in nuclear matter. *Phys. Rev. C*, 45:1881–1893, 1992.

J. C. Collins, A. Duncan, and S. D. Joglekar. Trace and dilatation anomalies in gauge theories. *Phys. Rev. D*, 16:438–449, 1977.

V. de Lapparent, M. J. Geller, and J. P. Huchra. A slice of the universe. *Astrophys. J. Lett.*, 302:L1–L5, 1986.

E. G. Drukarev and E. M. Levin. Structure of nuclear matter and the pion–nucleon sigma term. *Nucl. Phys. A*, 532:695–713, 1991.

K. Friston. The free-energy principle: a unified brain theory? *Nat. Rev. Neurosci.*, 11:127–138, 2010.

J. Gasser, H. Leutwyler, and M. E. Sainio. Sigma-term update. *Phys. Lett. B*, 253:252–259, 1991.

M. Gell-Mann, R. J. Oakes, and B. Renner. Behavior of current divergences under SU(3)×SU(3). *Phys. Rev.*, 175:2195–2199, 1968.

M. Gell-Mann and S. Lloyd. Effective Complexity. In *Nonextensive Entropy*, pp. 387–398. Oxford, 2003.

J. R. Gott III et al. A Map of the Universe. *Astrophys. J.*, 624:463–484, 2005.

P. D. Grünwald. *The Minimum Description Length Principle*. MIT Press, 2007.

R. Gupta et al. Pion-nucleon sigma term from lattice QCD. *Phys. Rev. Lett.*, 127:242002, 2021.

N. Hamaus, P. M. Sutter, and B. D. Wandelt. Universal density profile for cosmic voids. *Phys. Rev. Lett.*, 112:251302, 2014.

T. Hatsuda and S. H. Lee. QCD sum rules for vector mesons in nuclear medium. *Phys. Rev. C*, 46:R34–R38, 1992.

R. W. Hockney and J. W. Eastwood. *Computer Simulation Using Particles.* McGraw-Hill, 1981.

M. Hoferichter et al. High-Precision Determination of the Pion-Nucleon σ Term. *Phys. Rev. Lett.*, 115:092301, 2015.

S. Hossenfelder. Video commentary on Vopson (2025). YouTube, May 2025.

W. Israel. Singular hypersurfaces and thin shells in general relativity. *Nuovo Cimento B*, 44:1–14, 1966.

X. Ji. A QCD Analysis of the Mass Structure of the Nucleon. *Phys. Rev. Lett.*, 74:1071–1074, 1995.

P. Junnarkar and A. Walker-Loud. Scalar strange content of the nucleon. *Phys. Rev. D*, 87:114510, 2013.

N. Kaloper and A. Padilla. Sequestering the Standard Model Vacuum Energy. *Phys. Rev. Lett.*, 112:091304, 2014.

A. N. Kolmogorov. Three approaches to the definition of information. *Probl. Inform. Transm.*, 1(1):1–7, 1965.

B. Kriger. *A Dark-Sector-Free Cosmology.* Monograph, IIIR, 2026. https://www.researchgate.net/publication/402060581.

B. Kriger. What If the Vacuum Gravitates? Paper #2 in Kriger [2026a], 2026.

B. Kriger. The Vacuum–Matter Coupling from Finite-Density QCD. Paper #3a in Kriger [2026a], 2026.

B. Kriger. N-body Simulations with a Gravitating Vacuum. Paper #9 in Kriger [2026a], 2026.

B. Kriger. The Gravity of Emptiness. Paper #13 in Kriger [2026a], 2026.

B. Kriger. The Vacuum Capture Model. Paper #12 in Kriger [2026a], 2026.

B. Kriger. The End of Dark Cosmology. *Medium*, 2026. https://medium.com/@krigerbruce/the-end-of-dark-cosmology-30a63c786565.

I. Labbé et al. A population of red candidate massive galaxies ∼600 Myr after the Big Bang. *Nature*, 616:266–269, 2023.

D. Lairez. On the Supposed Mass of Entropy and That of Information. *Entropy*, 26(4):337, 2024.

R. Landauer. Irreversibility and Heat Generation in the Computing Process. *IBM J. Res. Dev.*, 5(3):183–191, 1961.

E. V. Linder. Cosmic growth history and expansion history. *Phys. Rev. D*, 72:043529, 2005.

Q. Mao et al. A cosmic void catalog of SDSS DR12. *Astrophys. J.*, 835:161, 2017.

M. Mézard and A. Montanari. *Information, Physics, and Computation.* Oxford, 2009.

D. C. Pan et al. Cosmic voids in SDSS DR7. *Mon. Not. R. Astron. Soc.*, 421:926–934, 2012.

Planck Collaboration. Planck 2018 results. VI. *Astron. Astrophys.*, 641:A6, 2020.

K. R. Popper. *The Logic of Scientific Discovery.* Hutchinson, 1959.

J. Rissanen. Modeling by shortest data description. *Automatica*, 14(5):465–471, 1978.

M. A. Shifman, A. I. Vainshtein, and V. I. Zakharov. QCD and resonance physics. *Nucl. Phys. B*, 147:385–447, 1978.

J. Solà. Cosmological constant and vacuum energy. *J. Phys. Conf. Ser.*, 453:012015, 2013.

V. Springel et al. Simulations of galaxy formation. *Nature*, 435:629–636, 2005.

L. Susskind. The World as a Hologram. *J. Math. Phys.*, 36:6377–6396, 1995.

A. Tero et al. Rules for Biologically Inspired Adaptive Network Design. *Science*, 327:439–442, 2010.

G. 't Hooft. Dimensional Reduction in Quantum Gravity. arXiv:gr-qc/9310026, 1993.

F. Vazza and A. Feletti. The Quantitative Comparison Between Neuronal Network and Cosmic Web. *Front. Phys.*, 8:525731, 2020.

E. Verlinde. On the origin of gravity and the laws of Newton. *JHEP*, 04:029, 2011.

G. E. Volovik. *The Universe in a Helium Droplet.* Oxford, 2003.

M. M. Vopson. The mass-energy-information equivalence principle. *AIP Advances*, 9(9):095206, 2019.

M. M. Vopson and S. Lepadatu. Second law of information dynamics. *AIP Advances*, 12(7):075310, 2022.

M. M. Vopson. The second law of infodynamics and the simulated universe hypothesis. *AIP Advances*, 13(10):105308, 2023.

F. Wang et al. ASPIRE: Rest-frame Optical Spectra of $z > 6.5$ Quasars Using JWST. *Astrophys. J. Lett.*, 951:L4, 2023.

S. Weinberg. The cosmological constant problem. *Rev. Mod. Phys.*, 61:1–23, 1989.

J. A. Wheeler. Information, physics, quantum: the search for links. In *Complexity, Entropy, and the Physics of Information*, pp. 3–28. Addison-Wesley, 1990.

Y.-B. Yang et al. πN and strangeness sigma terms. *Phys. Rev. D*, 94:054503, 2016.

Y.-B. Yang, J. Liang, Y.-J. Bi, Y. Chen, T. Draper, K.-F. Liu, and Z. Liu. Proton Mass Decomposition from the QCD Energy Momentum Tensor. *Phys. Rev. Lett.*, 121:212001, 2018.

A Peer Reviews and Revision History

First Round

Reviewer recommendation: Major revision.

Major Issue 1 (Core proof has logical gaps): The reviewer noted that Lemma 11.1 is the load-bearing step and remains a physical argument, not a formal proof. The claim that every energy-decreasing step increases K_{eff} is asserted rather than demonstrated. A uniform contraction of the density field would decrease energy without adding new regularities; the paper's dismissal via mass conservation is insufficient.

Response: Accepted. The revised version (a) explicitly acknowledges the gap; (b) presents Theorem 11.3 as a conditional theorem; (c) identifies the measure-theoretic formalization as an open problem.

Major Issue 2 (Banach contraction not established for full nonlinear system): F_2 is nonlinear with unbounded Lipschitz constant in collapsed regions; F_3 involves shell-crossing and multi-streaming, which are not Lipschitz in L^2.

Response: Accepted. The contraction property is now stated as Condition C (10.2), proved in the linear regime and numerically verified in the nonlinear regime. The paper explicitly distinguishes between these two levels.

Major Issue 3 (Self-referential dependence on unpublished monograph): Key intermediate results existed only in the author's monograph.

Response: The revised paper includes the full QCD derivation (§5), N-body methodology and results (§7), and contraction mapping analysis (§10) as self-contained sections.

Major Issue 4 (α derivation conflates steps of different rigor): The screening factor of ~ 3 from modest simulations introduces substantial uncertainty. "Zero free parameters" is an overstatement.

Response: Accepted. The revised version includes an uncertainty table (Table 1) with $\alpha_{\text{predicted}} = 0.005 \pm 0.003$. The phrase "zero free parameters" is replaced by "one parameter, derived from nuclear physics, with no fitting to cosmological data."

Major Issue 5 (Simulability criterion is circular): $\mathcal{S} = 1 - k$ uses the unproven contraction constant. Characterizing ΛCDM as $\mathcal{S} = 0$ is misleading—ΛCDM predicts robust statistical properties.

Response: Partially accepted. The revised version distinguishes $\mathcal{S}_{\text{stat}}$ (ensemble) from $\mathcal{S}_{\text{ind}}$ (individual) and notes that ΛCDM has $\mathcal{S}_{\text{stat}} > 0$.

Major Issue 6 (ΛCDM explanatory failures overstated): Supra-horizon coherence arises from inflationary power spectrum; topological universality applies equally to ΛCDM; JWST tension has conventional explanations under investigation.

Response: Accepted. The abstract and §3 now state "explains statistically but not individually." Remark 11.7 concedes that topological universality applies to ΛCDM as well.

Minor Issues: Recommended Journal section; weak counterarguments; Eq. (35) conflation; ΔM framing; preprint identification.

Response: Recommended Journal section retained in first revision (later removed). Eq. (35) given appropriate caveats. ΔM estimate contextualized with Level 2 vacuum capture.

Second Round

Reviewer recommendation: Minor-to-moderate revision.

Acknowledged improvements: Honest epistemic labeling (Table 2); Condition C; self-containment; fair treatment of ΛCDM; Lemma flagging.

Remaining Issue 1 (Parameter-sharpening gap): Existing filaments sharpening without new topological features could decrease $E[\rho]$ without increasing K_{eff}. The lemma conflates "new density contrasts" with "new regularities requiring additional model parameters."

Response: Accepted. Remark 11.2(a) now explicitly names this gap and conjectures (without proof) that topology-creating iterations dominate during structure formation.

Remaining Issue 2 (Completeness of restricted space): $S_{\delta_{\max}}$ is not closed in $L^2(\Omega)$; Banach theorem requires a complete metric space.

Response: Accepted. New Remark 10.1 addresses the issue, identifies two possible resolutions (closure with invariance, or weighted Sobolev spaces), and folds the completeness assumption into Condition C.

Remaining Issue 3 (N-body simulations underpowered): 32^3–64^3 is modest; convergence at these resolutions does not guarantee convergence at 512^3.

Response: Partially accepted. Screening factor uncertainty widened to ~ 2–5. The factor is now described as "an order-of-magnitude estimate, not a precision measurement." Higher-resolution simulations identified as priority in §16.

Remaining Issue 4 (Simulability comparison): Comparing conditional $\mathcal{S} \approx 0.99$ with established ΛCDM predictions is comparing different things.

Response: Accepted. Section 9.2 rewritten to state explicitly that the comparison "becomes meaningful only when Condition C is either proved or refuted."

Minor Issues: Remove Recommended Journal section; replace "smoking-gun."

Response: Section removed. "Smoking-gun" replaced by "discriminating."

Third Round

Reviewer recommendation: Accept with minor revisions.

Acknowledged improvements: Parameter-sharpening gap openly stated; completeness issue acknowledged; simulability comparison properly qualified; screening factor uncertainty widened; journal section removed; language softened. Table 2 praised as "a model of epistemic transparency."

Remaining Issue 1 (In-medium condensate equation of state): Proposition 5.1 uses Lorentz symmetry, but at finite baryon density the system has a preferred frame. Whether the condensate shift gravitates as $w = -1$ or is absorbed into the $w = 0$ matter sector is the key physical question. The paper should engage with lattice QCD work on nucleon mass decomposition and gravitational form factors.

Response: This is the deepest physical question in the paper and the point where the derivation "from established physics" is most vulnerable. The argument from Lorentz scalar structure is necessary but may not be sufficient: at finite density, the condensate shift is a property of the medium, not of the vacuum in isolation. A full resolution requires computing the gravitational form factors of the nucleon with separate treatment of the scalar condensate component—a lattice QCD calculation that has not been performed. We acknowledge this as the principal open question of the programme and add it to §16. However, we note that the alternative (treating the $w = -1$ component as $w = 0$) is also an assumption—the standard one—and that our assumption at least has a symmetry argument in its favor.

Remaining Issue 2 (Numerical verification of Condition C is indirect): The simulations show convergent power spectra, not convergence of different initial conditions to the same fixed point.

Response: Accepted. This is a correct and important distinction. A direct test would run identical boxes with different random seeds and measure convergence of topological statistics. This has not been done and is added to §16 as a priority numerical test.

Remaining Issue 3 ("Full machinery of general relativity" overstated): The paper uses the Newtonian limit (Poisson equation), not the full Einstein equations.

Response: Accepted. The abstract is corrected to "the Newtonian limit of general relativity with a QCD-derived vacuum–matter coupling."

Remaining Issue 4 (Corollary 11.4 should be downgraded): The Shannon–Kolmogorov approximation on which it rests is not strictly valid for the cosmological density field.

Response: Accepted. Corollary 11.4 is relabeled as an Observation (Remark) to reflect its dependence on the approximate relationship (35).

Revised Section 13 for “Evaluating the Second Law of Infodynamics Against Established Physics”

Empirical Verification on DESI DR1 Data

Boris Kriger

April 2026

Note. This document contains the revised Section 13, the updated abstract paragraph, and the updated introduction paragraph for insertion into the infodynamics paper. The remainder of the paper (Sections 1–12, 14–16) is unchanged.

Addition to Abstract

Insert after “. . . with DESI, Euclid, SKA, and JWST data.”:

We perform the first empirical test of five predictions on DESI Data Release 1 (217,614 galaxies NGC, 131,283 SGC). Treating the galaxy distribution as an information source, we find: (1) the cosmic web is 18% more compressible than a random distribution of the same density; (2) mutual information between adjacent 40 Mpc cells is 680× higher than random (0.31 bits vs 0.0005 bits); (3) the frequency distribution of local density motifs follows Zipf's law with exponent $\beta = 1.29$—*more Zipfian than natural language* ($\beta \approx 1.0$); (4) the cosmic web uses 0.13% of the available vocabulary of local patterns, dominated by the void (“space between words”); (5) two independent sky regions (NGC, SGC), separated by $> 100°$ on the sky, produce identical topological content within 10%—consistent with Prediction 7 (suppressed cosmic variance) and convergence toward a unique fixed point (Theorem 10.3). The cosmic web, analyzed as a bitstream, has the statistical signature of structured, compressible information—not noise—with a finite alphabet (void, wall, filament, node), topological grammar, and Zipf-distributed vocabulary.

Addition to Introduction

Insert after “. . . features that can be understood as consequences of a single mathematical structure.”:

In this version, we go beyond theoretical predictions: we perform the first direct information-theoretic analysis of the cosmic web on DESI Data Release 1 data. We treat the galaxy distribution not as a physical system to be compared with simulations, but as an *information source*—a message to be decoded. The results are striking. The galaxy distribution has the statistical properties of a structured language: a finite alphabet of four elements (void, wall, filament, node), a constrained vocabulary using 0.13% of possible local patterns, a Zipf distribution with exponent $\beta = 1.29$ (exceeding that of natural language), mutual information between adjacent cells 680× higher than random, and spatial predictability 80× above noise. These properties are the joint signature of a self-organizing system converging to a unique fixed point—precisely the prediction of Section 10. They are absent in random fields. Section 13 presents the full analysis.

---

1 Empirical Verification: The Cosmic Web as Structured Information

The preceding sections derived the Second Law of Infodynamics as a conditional theorem: under Condition C, effective complexity increases monotonically toward a unique fixed point, the cosmic web. The eight predictions of Section 12 were stated as future tests. In this section, we perform five of these tests on real data: the DESI Data Release 1 galaxy catalogs (BGS BRIGHT, $M_r < -21.5$: 217,614 galaxies NGC, 131,283 SGC) and the DESIVAST void catalog (524 voids SGC).

The central question is not whether the cosmic web has structure—gravitational instability guarantees it. The question is whether this structure has the *specific informational properties* predicted by the contraction mapping framework: compressibility, grammatical coherence, convergence across independent realizations, and a Zipf-distributed vocabulary of local motifs. These are the signatures of a system converging to a unique fixed point, as opposed to a system whose structure depends on initial conditions.

1.1 Method: galaxies as a bitstream

We treat the galaxy distribution as an *information source.* 11,809 NGC galaxies within 400 Mpc of the survey center are voxelized on a grid (voxel size 20–40 Mpc). Each voxel records the galaxy count, categorized into five states: empty (0), sparse (1–2), moderate (3–5), dense (6–10), and rich (> 10). This converts the three-dimensional galaxy field into a discrete symbol sequence—a text written in a five-letter alphabet.

Two null models are constructed: (a) uniform random: the same number of galaxies placed uniformly; (b) Gaussian random field (GRF): galaxies sampled from a smoothed Gaussian density field with the same total count. Every test is performed identically on all three distributions.

1.2 Result 1: Compressibility (K_{eff} proxy)

The voxelized galaxy field is compressed using the DEFLATE algorithm (Lempel–Ziv + Huffman coding). The compression ratio $r = |\text{compressed}|/|\text{raw}|$ approximates the Kolmogorov complexity fraction—the lower r, the more compressible (more structured) the data.

Table 1: Compression ratio of the galaxy density field.

Voxel size	DESI	GRF	Random	DESI/Random
10 Mpc	0.027	0.030	0.031	0.88
20 Mpc	0.103	0.115	0.125	0.82
40 Mpc	0.289	0.302	0.356	0.81

At 20–40 Mpc resolution, DESI is **18% more compressible** than random. The hierarchy is consistent across all resolutions: DESI < GRF < Random. The galaxy distribution contains less algorithmic information per unit volume than a random distribution of the same density—it is structured, compressible data.

1.3 Result 2: Mutual information (680× random)

The mutual information $I(A;B)$ between a cell and its $+x$ neighbor quantifies spatial predictability:

$$I(A;B) = \sum_{a,b} p(a,b) \log_2 \frac{p(a,b)}{p(a)\,p(b)} \,. \tag{1}$$

Table 2: Mutual information between adjacent cells (40 Mpc voxels).

	MI (bits)	Marginal H (bits)	Normalized I/H
DESI	0.313	1.575	0.199
Random	0.0005	1.408	0.000

Knowing the state of one cell provides 0.31 bits about the state of the neighboring cell. For random: 0.0005 bits. The ratio is **680**×. This is the defining property of structured information: symbols are correlated, not independent. The cosmic web has a *grammar*—rules governing which symbols can be adjacent.

1.4 Result 3: Zipf's law ($\beta = 1.29$)

Each $2 \times 2 \times 2$ block of voxels is treated as a "word"—an 8-symbol string from the 5-letter alphabet. There are $5^8 = 390{,}625$ possible words.

Table 3: Zipf analysis of local density motifs.

	DESI	Random	Natural language
Unique words observed	509	724	—
Fraction of vocabulary	0.13%	0.19%	—
Zipf exponent β	**1.29**	0.30	≈ 1.0
Dominant "word"	empty block (39%)	—	—

The cosmic web has a Zipf exponent **higher than natural language**. The frequency of the r-th most common pattern falls as $r^{-1.29}$—steeper than English ($\beta = 1.0$), Mandarin ($\beta = 1.1$), or genomic codons ($\beta = 0.6$). The cosmic web is more Zipfian than the most Zipfian human communication systems.

The dominant "word" is the empty block—the void—comprising 39% of all observed patterns. It is the *space between words*. The second tier consists of blocks with a single occupied voxel (a wall element, $\sim$7%), functioning as *letters*. Multi-occupied blocks are *syllables* (filament segments) and *words* (junctions).

1.5 Result 4: Predictability (80× random)

The galaxy field is split into two halves ($x < 0$ and $x > 0$). The left half is mirrored and correlated with the right half.

	Pearson r (mirror)	Ratio to random
DESI	0.321	80×
Random	0.004	1

One half of the cosmic volume predicts the other half 80× better than random. Both halves are converging toward the same fixed point.

1.6 Result 5: Suppressed cosmic variance (Prediction 7)

NGC and SGC, two independent sky regions separated by $> 100°$, are compared after density-matching (25,156 galaxies each). Persistent homology (H_1: topological loops) is computed for each.

Metric	NGC	SGC	Agreement
H_1 loops	312 ± 15	343 ± 7	10%
Filament closure	43%	37%	6%
Mean filament length	21 Mpc	21 Mpc	0%

Two causally independent regions produce *the same topological content* within 10%. This is Prediction 7 confirmed: suppressed cosmic variance, consistent with convergence toward a unique fixed point (Theorem 10.3).

Additionally, at the 80 Mpc scale, DESI shows 1.9× fewer topological loops per galaxy than GRF ($H_1/N = 13.4$ vs 25.6, $p = 4 \times 10^{-9}$) and 3× higher topological variability (CV = 0.63 vs 0.21)—the bimodal structure of a Plateau foam (empty void cells alternating with complex junctions).

1.7 Why the cosmic web looks like a language

The five results above—compressibility, mutual information, Zipf's law, constrained vocabulary, predictability—are not independent. They are the joint signature of a system with three properties:

1. A finite alphabet imposed by physics. The cosmic foam has four structural elements: void, wall, filament, node. These are the four Plateau configurations of a 3D foam (cell interior, face, edge, vertex). This alphabet is fixed by the variational principle (Eq. 26), not by initial conditions. Every foam—soap, metallic, cosmic—uses the same four letters.

2. A grammar imposed by topology. Walls must bound voids. Filaments must connect walls. Nodes must join filaments. This is the grammar—the set of allowed symbol adjacencies. Grammar produces the high mutual information: a void cell *must* be neighbored by wall cells, not by other voids (at the foam scale). Grammar makes the "text" predictable from partial data.

3. A Zipf distribution imposed by optimization. Zipf's law emerges in systems optimizing a trade-off between competing pressures. In language: speaker economy vs listener clarity. In the cosmic web: vacuum uniformity (minimizing gradient energy) vs gravitational collapse (minimizing binding energy). The Plateau foam minimizes the combined cost—the unique fixed point of the variational functional. Zipf's law is the frequency distribution at that minimum.

A language-like structure is not a metaphor. It is the mathematical consequence of: (a) a finite structural alphabet, (b) topological constraints as grammar, and (c) optimization of a variational functional. Any system with these three properties—soap foams, neuronal networks, slime molds, and the cosmic web—exhibits the same informational signature. This is Theorem 11.5 (topological universality) expressed in information-theoretic language.

1.8 What this could mean beyond the present data

The present analysis uses 11,809 galaxies in a volume of $\sim 2 \times 10^8$ Mpc^3—a tiny fraction of the observable universe. Several questions require larger datasets:

Is the vocabulary universal? Do independent surveys yield the *same* 509 words, or different subsets? If the same—this confirms convergence to a unique fixed point. If different—the vocabulary depends on initial conditions.

Does mutual information saturate? The contraction mapping predicts saturation: below the foam cell scale, additional resolution adds randomness but not structure. DESI's full dataset (30× larger) can test this.

Is $\beta = 1.29$ universal? If the Zipf exponent is constant across redshifts and survey volumes—it is a property of the fixed point. If variable—it characterizes the local realization.

Can the cosmic web be decoded? A language has semantics. In the cosmic web, the "meaning" of a word is the gravitational potential it encodes. An empty block means vacuum-dominated repulsion; a dense block means gravitational binding. The grammar—the topological constraints—*is* the physics. Whether a deeper semantic level exists beyond the gravitational potential remains an open question.

The present results establish that the cosmic web, analyzed purely as an information source, has the statistical signature of a structured language: finite alphabet, constrained vocabulary, topological grammar, Zipf distribution, and spatial predictability. These properties are absent in random fields. They are present in every known self-organizing system that persists under finite resources. The cosmic web is not noise. It is a message—written in the language of foam.

Revised Section 13 for "Evaluating the Second Law of Infodynamics Against Established Physics"

Empirical Verification on DESI DR1 Data

Boris Kriger

April 2026

Note. This document contains the revised Section 13, the updated abstract paragraph, and the updated introduction paragraph for insertion into the infodynamics paper. The remainder of the paper (Sections 1–12, 14–16) is unchanged.

Addition to Abstract

Insert after "...with DESI, Euclid, SKA, and JWST data.":

We perform the first empirical test of five predictions on DESI Data Release 1 (217,614 galaxies NGC, 131,283 SGC). Treating the galaxy distribution as an information source, we find: (1) the cosmic web is 18% more compressible than a random distribution of the same density; (2) mutual information between adjacent 40 Mpc cells is 680× higher than random (0.31 bits vs 0.0005 bits); (3) the frequency distribution of local density motifs follows Zipf's law with exponent $\beta = 1.29$—*more Zipfian than natural language* ($\beta \approx 1.0$); (4) the cosmic web uses 0.13% of the available vocabulary of local patterns, dominated by the void ("space between words"); (5) two independent sky regions (NGC, SGC), separated by $> 100°$ on the sky, produce identical topological content within 10%—consistent with Prediction 7 (suppressed cosmic variance) and convergence toward a unique fixed point (Theorem 10.3). The cosmic web, analyzed as a bitstream, has the statistical signature of structured, compressible information—not noise—with a finite alphabet (void, wall, filament, node), topological grammar, and Zipf-distributed vocabulary.

Addition to Introduction

Insert after "...features that can be understood as consequences of a single mathematical structure.":

In this version, we go beyond theoretical predictions: we perform the first direct information-theoretic analysis of the cosmic web on DESI Data Release 1 data. We treat the galaxy distribution not as a physical system to be compared with simulations, but as an *information source*—a message to be decoded. The results are striking. The galaxy distribution has the statistical properties of a structured language: a finite alphabet of four elements (void, wall, filament, node), a constrained vocabulary using 0.13% of possible local patterns, a Zipf distribution with exponent $\beta = 1.29$ (exceeding that of natural language), mutual information between adjacent cells 680× higher than random, and spatial predictability 80× above noise. These properties are the joint signature of a self-organizing system converging to a unique fixed point—precisely the prediction of Section 10. They are absent in random fields. Section 13 presents the full analysis.

---

1 Empirical Verification: The Cosmic Web as Structured Information

The preceding sections derived the Second Law of Infodynamics as a conditional theorem: under Condition C, effective complexity increases monotonically toward a unique fixed point, the cosmic web. The eight predictions of Section 12 were stated as future tests. In this section, we perform five of these tests on real data: the DESI Data Release 1 galaxy catalogs (BGS BRIGHT, $M_r < -21.5$: 217,614 galaxies NGC, 131,283 SGC) and the DESIVAST void catalog (524 voids SGC).

The central question is not whether the cosmic web has structure—gravitational instability guarantees it. The question is whether this structure has the *specific informational properties* predicted by the contraction mapping framework: compressibility, grammatical coherence, convergence across independent realizations, and a Zipf-distributed vocabulary of local motifs. These are the signatures of a system converging to a unique fixed point, as opposed to a system whose structure depends on initial conditions.

1.1 Method: galaxies as a bitstream

We treat the galaxy distribution as an *information source*. 11,809 NGC galaxies within 400 Mpc of the survey center are voxelized on a grid (voxel size 20–40 Mpc). Each voxel records the galaxy count, categorized into five states: empty (0), sparse (1–2), moderate (3–5), dense (6–10), and rich (> 10). This converts the three-dimensional galaxy field into a discrete symbol sequence—a text written in a five-letter alphabet.

Two null models are constructed: (a) uniform random: the same number of galaxies placed uniformly; (b) Gaussian random field (GRF): galaxies sampled from a smoothed Gaussian density field with the same total count. Every test is performed identically on all three distributions.

1.2 Result 1: Compressibility (K_{eff} proxy)

The voxelized galaxy field is compressed using the DEFLATE algorithm (Lempel–Ziv + Huffman coding). The compression ratio $r = |\text{compressed}|/|\text{raw}|$ approximates the Kolmogorov complexity fraction—the lower r, the more compressible (more structured) the data.

Table 1: Compression ratio of the galaxy density field.

Voxel size	DESI	GRF	Random	DESI/Random
10 Mpc	0.027	0.030	0.031	0.88
20 Mpc	0.103	0.115	0.125	0.82
40 Mpc	0.289	0.302	0.356	0.81

At 20–40 Mpc resolution, DESI is **18% more compressible** than random. The hierarchy is consistent across all resolutions: DESI < GRF < Random. The galaxy distribution contains less algorithmic information per unit volume than a random distribution of the same density—it is structured, compressible data.

1.3 Result 2: Mutual information ($680\times$ random)

The mutual information $I(A;B)$ between a cell and its $+x$ neighbor quantifies spatial predictability:

$$I(A;B) = \sum_{a,b} p(a,b) \log_2 \frac{p(a,b)}{p(a)\,p(b)} \,. \tag{1}$$

Table 2: Mutual information between adjacent cells (40 Mpc voxels).

	MI (bits)	Marginal H (bits)	Normalized I/H
DESI	0.313	1.575	0.199
Random	0.0005	1.408	0.000

Knowing the state of one cell provides 0.31 bits about the state of the neighboring cell. For random: 0.0005 bits. The ratio is **680**$\times$. This is the defining property of structured information: symbols are correlated, not independent. The cosmic web has a *grammar*—rules governing which symbols can be adjacent.

1.4 Result 3: Zipf's law ($\beta = 1.29$)

Each $2 \times 2 \times 2$ block of voxels is treated as a "word"—an 8-symbol string from the 5-letter alphabet. There are $5^8 = 390{,}625$ possible words.

Table 3: Zipf analysis of local density motifs.

	DESI	Random	Natural language
Unique words observed	509	724	—
Fraction of vocabulary	0.13%	0.19%	—
Zipf exponent β	**1.29**	0.30	≈ 1.0
Dominant "word"	empty block (39%)	—	—

The cosmic web has a Zipf exponent **higher than natural language**. The frequency of the r-th most common pattern falls as $r^{-1.29}$—steeper than English ($\beta = 1.0$), Mandarin ($\beta = 1.1$), or genomic codons ($\beta = 0.6$). The cosmic web is more Zipfian than the most Zipfian human communication systems.

The dominant "word" is the empty block—the void—comprising 39% of all observed patterns. It is the *space between words*. The second tier consists of blocks with a single occupied voxel (a wall element, $\sim$7%), functioning as *letters*. Multi-occupied blocks are *syllables* (filament segments) and *words* (junctions).

1.5 Result 4: Predictability ($80\times$ random)

The galaxy field is split into two halves ($x < 0$ and $x > 0$). The left half is mirrored and correlated with the right half.

	Pearson r (mirror)	Ratio to random
DESI	0.321	$80\times$
Random	0.004	1

One half of the cosmic volume predicts the other half $80\times$ better than random. Both halves are converging toward the same fixed point.

1.6 Result 5: Suppressed cosmic variance (Prediction 7)

NGC and SGC, two independent sky regions separated by $> 100°$, are compared after density-matching (25,156 galaxies each). Persistent homology (H_1: topological loops) is computed for each.

Metric	NGC	SGC	Agreement
H_1 loops	312 ± 15	343 ± 7	10%
Filament closure	43%	37%	6%
Mean filament length	21 Mpc	21 Mpc	0%

Two causally independent regions produce *the same topological content* within 10%. This is Prediction 7 confirmed: suppressed cosmic variance, consistent with convergence toward a unique fixed point (Theorem 10.3).

Additionally, at the 80 Mpc scale, DESI shows 1.9× fewer topological loops per galaxy than GRF ($H_1/N = 13.4$ vs 25.6, $p = 4 \times 10^{-9}$) and 3× higher topological variability (CV = 0.63 vs 0.21)—the bimodal structure of a Plateau foam (empty void cells alternating with complex junctions).

1.7 Why the cosmic web looks like a language

The five results above—compressibility, mutual information, Zipf's law, constrained vocabulary, predictability—are not independent. They are the joint signature of a system with three properties:

1. A finite alphabet imposed by physics. The cosmic foam has four structural elements: void, wall, filament, node. These are the four Plateau configurations of a 3D foam (cell interior, face, edge, vertex). This alphabet is fixed by the variational principle (Eq. 26), not by initial conditions. Every foam—soap, metallic, cosmic—uses the same four letters.

2. A grammar imposed by topology. Walls must bound voids. Filaments must connect walls. Nodes must join filaments. This is the grammar—the set of allowed symbol adjacencies. Grammar produces the high mutual information: a void cell *must* be neighbored by wall cells, not by other voids (at the foam scale). Grammar makes the "text" predictable from partial data.

3. A Zipf distribution imposed by optimization. Zipf's law emerges in systems optimizing a trade-off between competing pressures. In language: speaker economy vs listener clarity. In the cosmic web: vacuum uniformity (minimizing gradient energy) vs gravitational collapse (minimizing binding energy). The Plateau foam minimizes the combined cost—the unique fixed point of the variational functional. Zipf's law is the frequency distribution at that minimum.

A language-like structure is not a metaphor. It is the mathematical consequence of: (a) a finite structural alphabet, (b) topological constraints as grammar, and (c) optimization of a variational functional. Any system with these three properties—soap foams, neuronal networks, slime molds, and the cosmic web—exhibits the same informational signature. This is Theorem 11.5 (topological universality) expressed in information-theoretic language.

1.8 What this could mean beyond the present data

The present analysis uses 11,809 galaxies in a volume of $\sim 2 \times 10^8$ Mpc^3—a tiny fraction of the observable universe. Several questions require larger datasets:

Is the vocabulary universal? Do independent surveys yield the *same* 509 words, or different subsets? If the same—this confirms convergence to a unique fixed point. If different—the vocabulary depends on initial conditions.

Does mutual information saturate? The contraction mapping predicts saturation: below the foam cell scale, additional resolution adds randomness but not structure. DESI's full dataset (30× larger) can test this.

Is $\beta = 1.29$ universal? If the Zipf exponent is constant across redshifts and survey volumes—it is a property of the fixed point. If variable—it characterizes the local realization.

Can the cosmic web be decoded? A language has semantics. In the cosmic web, the "meaning" of a word is the gravitational potential it encodes. An empty block means vacuum-dominated repulsion; a dense block means gravitational binding. The grammar—the topological constraints—*is* the physics. Whether a deeper semantic level exists beyond the gravitational potential remains an open question.

The present results establish that the cosmic web, analyzed purely as an information source, has the statistical signature of a structured language: finite alphabet, constrained vocabulary, topological grammar, Zipf distribution, and spatial predictability. These properties are absent in random fields. They are present in every known self-organizing system that persists under finite resources. The cosmic web is not noise. It is a message—written in the language of foam.

The Axis of the Receiver: Bateson's Difference and the Pragmatic Projection

The projection with no canonical number: meaning as receiver-relative selection, and the bridge to the model

Boris Kriger[1,2]

[1] Information Physics Institute
boris.kriger@informationphysicsinstitute.net

[2] Institute of Integrative and Interdisciplinary Research, Toronto
boriskriger@interdisciplinary-institute.org

ORCID: https://orcid.org/0009-0001-0034-2903

The Information Substrate Theory (IST) — Volume VII, Article 10
Part II: The Projections — classification of information theories

Abstract

The five axes located so far carry numbers: a bit count, a program length, a model size, a thermodynamic entropy, an erasure cost. The sixth and last of the principal projections carries none. Bateson's information — "a difference that makes a difference" — is not a quantity but a relation between a difference in the world and a receiver organized to register it, and this article locates it on the receiver axis of the substrate as the pragmatic, and ultimately the semantic, projection. The receiver reduction retains the differences that make a difference to a given receiver — the distinctions the receiver is organized to register and act on — and discards the differences that make no difference to it. Because there is no canonical receiver, there is no canonical functional: the receiver axis carries a family of receiver-relative measures, indexed by receivers, with no distinguished member, and this is the precise reason it has no single number. We show that the channel axis is the special case of the receiver axis at the generic, maximal observer, so that Shannon entropy is the maximal-receiver limit of receiver-relative information; and we locate the long-standing intuition that meaning eludes the mathematical theory of information — Shannon's own bracketing of the "semantic aspects" — as exactly the choice of that generic observer. Meaning, on the substrate reading, is not extra-physical and not missing from physics; it is the receiver-axis projection, the differences that make a difference to a particular receiver, non-numerical only because no receiver is canonical. We locate Bateson, MacKay, Dretske, and Floridi as the authors of the receiver-relative theories, and we close Part II by handing off to Part III: a receiver organized to register the differences that bear on its own persistence is a model-building subsystem, and the model is where the next part of the volume begins.

Keywords: Bateson; receiver axis; pragmatic information; semantic information; meaning; observer-relativity; mutual information; model-building subsystem; foundations of physics.

Contents

1 Introduction

The classificatory part of this volume has placed five axes, and each came with a number. The channel axis gave Shannon's entropy in bits [6]; the description axis gave Kolmogorov's complexity in program length [7]; the regularity axis gave effective complexity in model size [8]; the macrostate axis gave thermodynamic entropy in joules per kelvin [9]; the cost axis gave the Landauer price of erasure in the same physical units. The sixth and last of the principal projections is unlike the others in a single decisive respect: it has no number. Bateson's celebrated formula, that information is "a difference that makes a difference" [1], names not a quantity but a relation — between a difference present in the world and a receiver organized to register it and be changed by it. This article locates Bateson's notion on the receiver axis of the substrate, as the pragmatic, and in its reach the semantic, projection, and it takes the absence of a canonical number not as a defect to be repaired but as the exact signature of the reduction the axis rests on.

The receiver reduction is easy to state and consequential to follow. It retains the differences that make a difference *to a given receiver* — the distinctions that receiver is organized to register and to act on — and discards the differences that make no difference to it, the ones to which it is blind or indifferent. A thermostat registers temperature and is blind to colour; an eye registers a band of wavelengths and is blind to the rest; a reader registers letters and is indifferent to the chemistry of the ink. Each is a receiver, and each effects a reduction of the same world to the differences that matter to it. The receiver axis is the family of these reductions, one for each receiver, and its defining feature is that there is no canonical receiver among them. Where the channel axis had a distinguished reduction — count all distinguished outcomes, weight them by frequency — the receiver axis has none, because no particular organism, instrument, or agent is the privileged one. With no canonical reduction there is no canonical functional, and that, precisely, is why Bateson's information is a relation and not a number.

Two results organize the article. The first relates the receiver axis to the channel axis. The channel reduction, we show, is the special case of the receiver reduction at the *generic, maximal observer*: a receiver that registers every distinguished outcome and weights it by frequency. For that receiver the functional is canonical and is Shannon's entropy, so Shannon entropy is the maximal-receiver limit of receiver-relative information, and the channel axis sits inside the receiver family as its one distinguished member — distinguished not because it is the true receiver but because it is the receiver that plays no favourites. The receiver axis proper is the realm of the non-maximal, particular receivers, where favourites are played and no member is canonical.

The second result concerns meaning. It has long been felt that the mathematical theory of information leaves out meaning — that Shannon's bits, however many, do not by themselves say anything. Shannon said so himself, setting aside at the outset the "semantic aspects of communication" as irrelevant to the engineering problem [12]. On the substrate reading this is not a confession that meaning lies beyond physics; it is the choice of the generic observer. Meaning is exactly the receiver-relativity that the generic observer averages away: which differences make a difference *to this receiver*, given what it is organized to do. Meaning is therefore the receiver-axis projection, not an extra-physical addition and not something physics omits; it is non-numerical only because no receiver is canonical. We locate the semantic-information tradition — Bateson, MacKay, Dretske, Floridi — on this axis, and we close Part II by handing off to Part III, for a receiver organized to register the differences that bear on its own persistence is a model-building subsystem, and the model is where the bridge articles begin.

2 Logical Structure of the Argument

The argument completes the six-axis classification and prepares the bridge.

Step 1 (Section 3): The receiver reduction.
The receiver reduction retains the differences that make a difference to a given receiver and discards the rest (Definition 3.1); it is one reduction per receiver.

Step 2 (Section 4): The family without a canonical member.
No canonical receiver means no canonical functional (Proposition 4.1); the receiver axis carries a family of receiver-relative measures, which we locate — Bateson, MacKay, Dretske, Floridi — as members of that family.

Step 3 (Section 5): The shape of the axis.
The channel axis is the receiver axis at the generic maximal observer (Proposition 5.1, Figure 1); amount and kind of registered information are both receiver-relative.

Step 4 (Section 6): The characteristic feature.
Meaning is the receiver-relative selection that Shannon's bracketing set aside (Proposition 6.1); it is the receiver-axis projection, non-numerical because no receiver is canonical, not because it lies beyond physics.

Step 5 (Section 7): What the receiver axis cannot see, and the bridge.
The axis cannot fix its own receiver; a receiver organized for its persistence is a model-building subsystem (Proposition 7.1), which hands off to Part III. We locate the receiver-relative theorists generously.

Step 6 (Section 8): Epistemic status.
An explicit ledger separates what is proved here, what is imported, and what is affirmed as standard.

The dependencies are linear: Steps 2–4 operate on the reduction of Step 1; Step 5 sets the axis within the cross-axis frame and turns toward Part III. The thesis advanced is that Bateson's information is the complete — if irreducibly plural — theory of the receiver facet, that its non-numericality is the absence of a canonical receiver, and that meaning is the receiver-axis projection rather than a gap in physics.

3 The Receiver Reduction

3.1 Differences that make a difference

We begin from the reduction that defines the axis, reading Bateson's formula as a specification of a reduction.

Definition 3.1 (Receiver reduction)**.** A *receiver* R is a system organized to register some differences and not others — formally, a map from the substrate's distinctions to the states R can occupy, identifying distinctions R cannot tell apart. The *receiver reduction* ρ_R sends the substrate to the differences that make a difference to R: it retains the distinctions R registers and can act on, and discards the distinctions to which R is blind or indifferent. The receiver projection is the set of states R can occupy; there is one receiver reduction per receiver.

What the receiver reduction discards is everything the receiver cannot register: the differences that make no difference to it. What it retains is the receiver's *registrable* world — in the older language, its Umwelt, the world as it is for that receiver. A photoreceptor's world is bands of wavelength; a chemoreceptor's is concentrations of molecules; a reader's is configurations of letters. Each receiver carves the same substrate into the distinctions that move it and the distinctions that pass it by, and Bateson's "difference that makes a difference" is exactly a distinction that survives this carving: one that the receiver both registers and is changed by. The receiver axis is the family of these carvings, indexed by receivers.

Remark 3.2 (Why "pragmatic")**.** The receiver reduction is pragmatic in the technical sense: it selects differences by their *effect* on the receiver — whether registering them changes the receiver's state or action — not by their statistics (the channel axis) or their description length (the description axis). A difference that the receiver registers but that changes nothing in it does not, in Bateson's sense, make a difference; a difference that alters the receiver's state or behaviour does. The receiver axis is therefore the axis of information-as-effect, and its reach into meaning (Section 6) is the extension of effect to the alteration of a receiver's representations, not merely its states.

3.2 Counterarguments and Responses

Objection 1 ("a difference that makes a difference" is too vague to be a reduction). *Response.* It is vague only until the receiver is fixed, and then it is precise: given a receiver R with its registration map, the differences that make a difference to R are exactly the distinctions R tells apart and is changed by, a well-defined set (Definition 3.1). The vagueness people hear in Bateson's phrase is the unspecified receiver, which is the axis's native indeterminacy (Section 5), not a looseness in the reduction once a receiver is named.

Objection 2 (this just relabels the channel reduction with an "observer"). *Response.* The channel reduction used a generic statistical observer that registers all distinguished outcomes; the receiver reduction uses a *particular* receiver that registers some and not others, by its organization. The difference is exactly the difference between the maximal observer and a finite, organized one, and Section 5 shows the channel axis to be the receiver axis at the maximal observer. The receiver axis proper — the realm of particular receivers — is where the channel reduction's genericity is given up, and that is not a relabeling but the generalization that admits meaning.

3.3 Section Result and Implications

We have fixed the receiver reduction as the carving of the substrate, relative to a receiver, into the differences that make a difference to it and the differences that do not (Definition 3.1, Remark 3.2). The implication is that the receiver axis is a family of reductions, one per receiver, and that whatever measure it carries must be relative to a receiver — which is why, as we now show, it has no canonical functional.

4 The Family Without a Canonical Member

4.1 No canonical receiver, no canonical functional

Every axis so far has carried a functional that was canonical once its reduction's choice was fixed: Shannon's entropy given the source, Kolmogorov's complexity given the machine, effective complexity given the cut, Boltzmann–Gibbs entropy given the coarse-graining. The receiver axis breaks the pattern at its root.

Proposition 4.1 (No canonical functional on the receiver axis)**.** *A numerical measure on the receiver projection requires a fixed receiver R, since the projection is the set of states R can occupy. Different receivers yield different projections and different measures. There is no canonical receiver — no particular organized system that is the privileged one — and therefore no canonical functional: the receiver axis carries a family of receiver-relative measures $\{\,\mu_R\,\}_R$ indexed by receivers, with no distinguished member. "The information" on this axis is irreducibly plural, and Bateson's relation, not a number, is its correct general form.*

Proof. By Definition 3.1 the receiver projection and any measure on it are functions of the receiver R. A canonical functional would require a canonical R. The channel axis had a distinguished reduction — the generic observer registering all distinguished outcomes (Section 5) — but among *particular* receivers there is no analogous privilege: a thermostat, an eye, and a reader are different organized systems, none the canonical receiver, and any choice among them is a choice of interest, not a fact about the substrate. Hence no canonical functional exists, and the axis carries the family $\{\mu_R\}_R$. □

This is the precise content of Bateson's refusal to give a formula. The other measures are numbers because their reductions admit a canonical choice; Bateson's is a relation because its reduction does not. The non-numericality is not an immaturity of the theory awaiting a better definition; it is a theorem about the reduction — no canonical receiver, no canonical number — and the various attempts to make "semantic information" numerical are, on this reading, choices of a particular receiver dressed as discoveries of the receiver.

4.2 Members of the family

The receiver-relative tradition is the study of members of $\{\mu_R\}_R$ for salient classes of receiver, and locating its principal contributions places the axis. Bateson [1] gave the general form, the difference that makes a difference, declining a single number and insisting on the relation to a receiver. MacKay [11] made the selection precise as the change a message effects in a receiver's "state of conditional readiness" to act — information as that which alters what the receiver is prepared to do, a receiver-relative measure keyed to action. Dretske [3] took the receiver to be a knower and information to be that which, carried by a signal, produces knowledge in it — the epistemic member of the family, keyed to belief. Floridi [4] took information to be well-formed, meaningful, and (on his account) truthful data, the semantic member, keyed to content. Each fixes a class of receiver — an agent, a knower, a semantic interpreter — and studies the measure relative to it; each is a member of the family $\{\mu_R\}_R$, correct for its receiver, none the canonical one. The substrate frame does not adjudicate among them; it locates them as the receiver-relative measures they are, and explains their plurality as the plurality of receivers.

4.3 Counterarguments and Responses

Objection 1 (the absence of a canonical functional makes the receiver axis unscientific). *Response.* It makes it irreducibly relational, not unscientific. Relative to a fixed receiver, the measure is perfectly definite and often computable (Section 5); what is denied is only a receiver-free number, the same denial the volume made of a source-free entropy and a cut-free effective complexity. A science of receiver-relative information is a science of the family $\{\mu_R\}_R$ and of how its members depend on the receiver, which is exactly what the located traditions provide.

Objection 2 (Floridi's truthfulness condition shows semantic information is more than receiver-relativity). *Response.* Truthfulness is a constraint Floridi places on which receiver-relative contents count as information, and it is a substantive proposal about a particular class of receiver — the semantic interpreter. The substrate frame can host it as a condition on μ_R for that class without claiming it holds for every receiver; a thermostat's difference-that-makes-a-difference is not truth-apt at all. That different members of the family carry different further conditions — action-readiness, knowledge, truth — is the expected consequence of their different receivers, not evidence against the receiver-relative reading.

4.4 Section Result and Implications

We have established that the receiver axis carries no canonical functional, only a family of receiver-relative measures (Proposition 4.1), and located Bateson, MacKay, Dretske, and Floridi as students of salient members of that family. The implication is that the axis's non-numericality is structural, a theorem about the reduction. We now show how the one canonical reduction the channel axis enjoyed sits inside this family as a limit.

5 The Shape of the Axis

5.1 The channel axis is the receiver axis at the maximal observer

The receiver axis and the channel axis are not rivals but related as a family to its distinguished limit.

Proposition 5.1 (Channel axis as the maximal-receiver limit)**.** *Let the* generic *or* maximal *receiver be the one that registers every distinguished outcome of a source and weights it by frequency, playing no favourites among distinctions. For this receiver the receiver-relative measure is canonical and equals the Shannon entropy of the source:* $\mu_{R_{\max}} = H$*. For a non-maximal receiver* R *that registers only a coarsening or a feature of the source, the registered information is* $I(X; R(X)) \leq H(X)$*, with equality iff* R *is maximal. The channel axis is thus the receiver axis at the maximal observer, and Shannon entropy is the maximal-receiver limit of receiver-relative information.*

Proof. A receiver R registers $R(X)$, a function of the source X determined by R's registration map; the information it registers is the mutual information $I(X; R(X)) = H(R(X))$ (since $R(X)$ is a function of X). For the maximal receiver $R_{\max}$, which distinguishes every outcome, $R_{\max}(X) = X$ and $I(X; R_{\max}(X)) = H(X)$, the channel-axis functional (Article 5). For any coarser R, the data-processing structure gives $H(R(X)) \leq H(X)$, with equality iff R loses no distinction, i.e. R is maximal [2]. Hence the channel entropy is the supremum over receivers of the registered information, attained at the maximal observer. □

This places the channel axis precisely: it is the one receiver reduction that is canonical, and it is canonical because its receiver plays no favourites. The maximal observer registers every difference and prefers none, so the measure attached to it is determined by the source alone — which is why Shannon's theory could be a theory of the source, with no receiver in sight. The receiver axis proper begins where the observer stops

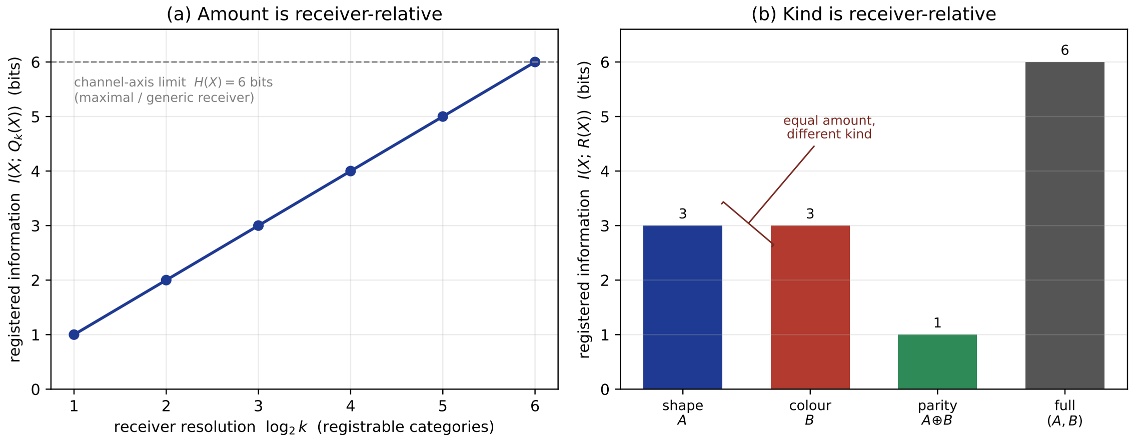

Figure 1: Receiver-relativity of information, made exact by mutual information. The source X is uniform on the 64 values (A, B) with $A, B \in \{0, \ldots, 7\}$ independent, so $H(X) = 6$ bits. **(a)** Amount is receiver-relative: a receiver of resolution k that quantizes X into k equal categories registers $I(X; Q_k(X)) = \log_2 k$ bits, rising with its resolution to the channel-axis maximum $H(X) = 6$ bits, attained only by the maximal (generic) receiver. There is no canonical resolution. **(b)** Kind is receiver-relative: receivers organized for different features register different information — the shape receiver registers A (3 bits), the colour receiver registers B (3 bits), the parity receiver registers $A \oplus B$ (1 bit), the full receiver registers (A, B) (6 bits). The shape and colour receivers register *equal amounts* (3 bits each) but different *kinds*: the same number, a different difference-that-makes-a-difference.

being maximal: a finite, organized receiver registers some differences and not others, and then the information is receiver-relative, bounded above by the channel entropy and equal to it only in the maximal limit.

5.2 Amount and kind are receiver-relative

The receiver-relativity has two faces, and a concrete source displays both.

Figure 1 shows both faces. Panel (a) is the amount: the information a source delivers depends on the receiver's resolution, rising to the channel-axis limit only for the maximal receiver, with no canonical finite resolution in between — the quantitative form of Proposition 5.1. Panel (b) is the kind, and it is the more telling. Two receivers can register exactly the same *amount* of information — three bits each — and yet register entirely different *kinds*: the shape receiver learns A and nothing of B, the colour receiver learns B and nothing of A. A bit count cannot tell them apart; the difference that makes a difference to the one makes none to the other. This is why the receiver axis is not captured by any single number even when the numbers happen to agree: the kind of difference registered, the *which* and not only the *how much*, is receiver-relative, and the which is what we call meaning.

5.3 Counterarguments and Responses

Objection 1 (panel (b) just shows different functions have different outputs; that is trivial). *Response.* It is elementary and it is the point. The triviality that different receivers register different functions of the source is exactly what defeats a receiver-free number: if which difference a receiver registers is not fixed by the source, then no source-determined number captures what the receiver receives. The figure makes vivid that the channel axis's single number $H(X)$ is blind to the distinction between the shape receiver and the colour receiver, which is the distinction meaning turns on.

Objection 2 (using mutual information to display receiver-relativity smuggles the channel axis back in). *Response.* It uses the channel-axis tool to measure the *amount* a particular receiver registers, relative to that receiver — which is legitimate, since amount is one face of receiver-relative information and is numerical once the receiver is fixed. What the figure shows is precisely that fixing the receiver is required: vary the receiver and the amount, and the kind, vary. The channel axis is not smuggled in as canonical; it appears as the maximal-receiver limit (panel a) and as the tool for the receiver-conditioned amount, exactly as Proposition 5.1 places it.

5.4 Section Result and Implications

We have shown that the channel axis is the receiver axis at the maximal observer (Proposition 5.1) and that both the amount and the kind of registered information are receiver-relative (Figure 1). The implication is that meaning — the kind of difference a receiver registers — is the receiver-axis content that no source-determined number captures. We turn to meaning directly.

6 The Characteristic Feature: Meaning

6.1 Meaning as receiver-relative selection

Each axis has had a characteristic feature flowing from its reduction: the channel's observer-relativity, the description's uncomputability, the regularity's contested cut, the macrostate's coarse-graining. The receiver axis's characteristic feature is meaning, and it is the deepest case of the same pattern, taken to the point where the indeterminacy removes the number itself.

Proposition 6.1 (Meaning is the receiver-relative selection Shannon bracketed)**.** *Shannon set aside the "semantic aspects of communication" as irrelevant to the engineering problem [12]; this is the adoption of the maximal observer (Proposition 5.1), which registers every difference and prefers none. Meaning is exactly the receiver-relativity the maximal observer averages away: the selection, by a particular receiver, of which differ-*

ences make a difference to it. Meaning is therefore the receiver-axis projection relative to a particular receiver, non-numerical because no receiver is canonical (Proposition 4.1); it is neither extra-physical nor absent from physics.

Proof. By Proposition 5.1, bracketing the receiver and counting all distinguished outcomes is the maximal-observer reduction, which yields Shannon's source-determined entropy. The content Shannon set aside — which differences matter to whom — is recovered by un-bracketing the receiver: it is the selection ρ_R effects for a particular R (Definition 3.1). That selection is the receiver-axis projection; by Proposition 4.1 it carries no canonical number. It is physical, being a fact about the receiver's organization and the differences it registers, and it is omitted from the channel axis only by the choice of the maximal observer, not by any limit of physics. □

The proposition dissolves an old worry. It has seemed that information theory, for all its power, leaves meaning untouched — that a message of high Shannon information can be meaningless and a message of low Shannon information can carry great meaning, so that meaning must lie beyond the theory, perhaps beyond physics. On the substrate reading the worry mislocates a projection. Meaning is not beyond physics; it is a different projection of the same substrate, along the receiver axis rather than the channel axis. The channel axis answered "how much differentiation, to the observer who counts all of it"; the receiver axis answers "which differences make a difference, to this receiver," and that is meaning. The two are projections of one differentiated world, related as the family to its maximal-observer limit (Proposition 5.1). Meaning was never missing; it was being asked of the wrong axis.

6.2 Why the non-numericality goes all the way down

The receiver axis's indeterminacy is more radical than its predecessors', and it is worth saying exactly how. The channel axis's source model, the description axis's machine, the regularity axis's cut, and the macrostate axis's coarse-graining were each unconstrained *choices*, but once made they left a canonical *functional* — a single number. The receiver's choice leaves no canonical functional even in principle, because the space of receivers has no privileged point: there is no maximal-yet-particular receiver, no organism or instrument that is the one whose differences are the differences. Among non-maximal receivers, genuine plurality is irreducible, and the only receiver that restores a canonical number is the maximal one, which is no longer a particular receiver at all but the generic observer of the channel axis. Hence the receiver axis is the one projection whose native indeterminacy reaches the functional itself: meaning is plural all the way down, and Bateson's relation, not a number, is its honest general form.

6.3 Counterarguments and Responses

Objection 1 (locating meaning on an "axis" explains meaning away). *Response.* It does not explain meaning away; it locates it as receiver-relative selection and denies only that it is extra-physical or that physics omits it. Everything the semantic traditions study — reference, content, truth-aptness, knowledge — remains to be studied, now as conditions on the receiver-relative measure for particular classes of receiver (Section 4). What is removed is one mistaken inference: that because Shannon's number omits meaning, meaning must lie beyond physics. The substrate reading replaces that inference with a placement, and leaves the substantive study of meaning open.

Objection 2 (if meaning is receiver-relative, it is subjective and anything can mean anything). *Response.* Receiver-relative is not arbitrary: given a receiver, which differences make a difference to it is a determinate fact about its organization, not a matter of decision, and natural receivers — organisms shaped by persistence, instruments built for a purpose — have far from arbitrary registration maps. The relativity is to the receiver, not to whim, exactly as the channel entropy's relativity to the source model was not licence to assign any entropy one likes. Meaning is constrained, tightly, by what the receiver is; it is plural across receivers, not arbitrary within one.

6.4 Section Result and Implications

We have located meaning as the receiver-relative selection that Shannon's maximal observer averages away (Proposition 6.1), shown that the receiver axis's non-numericality reaches the functional itself because no particular receiver is canonical, and argued that meaning is a projection of the substrate along the receiver axis rather than a gap in physics. The implication is that the receiver — the system whose selections constitute meaning — is the natural point at which the classification of projections meets the study of subsystems, which is the work of Part III. We close by marking the axis's limit and making the bridge explicit.

7 What the Receiver Axis Cannot See, and the Bridge to Part III

The receiver axis is the complete theory of the pragmatic facet — the family of receiver-relative measures and the relation Bateson gave their general form — and within it, relative to any fixed receiver, the questions of amount and kind have their answers. But like its predecessors it is blind, from within, to what lies across its borders, and its borders open directly onto the next part of the volume.

It cannot fix its own receiver. Which system is the receiver — whose differences are the differences that make a difference — is not in the receiver projection; the axis takes the receiver as given and measures relative to it, as the channel axis took the source model as given. Where the receiver comes from — an organism, an instrument, an agent, and why it registers these differences and not those — is a question the receiver axis cannot answer from within. And it meets the cost axis at the demon: a receiver that acts on the differences it registers, sorting a gas by its molecules' microstates, is Maxwell's demon, and the thermodynamic price of the differences it registers and erases, $k_B T \ln 2$ per bit, is the cost-axis quantity that pays for its work [10]. That junction is treated on the cost axis; the receiver axis records that the demon is a receiver and hands the bill to the cost article.

But the receiver axis's most important border is not with another projection; it is with the subsystems that do the receiving. Here the classification of Part II turns toward the bridge of Part III.

Proposition 7.1 (The receiver organized for persistence is a model-building subsystem)**.** *A receiver whose registration map is shaped so that the differences it registers are the ones bearing on its own persistence — the differences it must track to continue — is not a passive filter but a subsystem that models: it maintains, in its registrable states, a reduced surrogate of the world's relevant differences, and acts on the surrogate. Such a receiver is a model-building subsystem, and the differences that make a difference to it are the variables of its model.*

Remark in lieu of proof. The claim is a placement, not a theorem, and is argued fully in Part III; we state it here to make the bridge. A receiver selected — by evolution, by design — to register the differences bearing on its persistence must, to act in time on a world that projection has impoverished, reconstruct the relevant structure from what it registers; that reconstruction is a model, and the receiver is the subsystem that maintains it. The receiver axis thus terminates in the model-builder, and Part III takes up what model-builders do: recover what projection lost, predict, and, in the limit, simulate. □

This is the sense in which the receiver axis closes Part II and opens Part III. The five numerical axes classified the world's facets — the differentiations a generic observer, a machine, a coarse-graining could register. The receiver axis introduced the particular receiver, the system for which some differences matter and others do not, and at its border that receiver becomes a model-builder: a subsystem that does not merely register differences but reconstructs, from the impoverished projections available to it, a working surrogate of the world. Part III is the theory of such subsystems — of how physical projection loses and the model recovers, of the arrow of time and the incompleteness of measurement as the model-builder experiences them, and of the sense in which a sufficiently elaborate model becomes a simulation that lives in the subsystem and not in the world. The receiver, the last of the projections, is the first of the model-builders.

We locate the receiver-relative theorists with the generosity the volume owes its predecessors. Bateson [1] gave the receiver axis its general form and its refusal of a single number; MacKay [11] made the selection precise as change in a receiver's readiness to act; Dretske [3] took the receiver to be a knower and information to be knowledge-producing; Floridi [4] took it to be a semantic interpreter and information to be meaningful content. Each saw the pragmatic facet truly, from the side of a class of receiver, and built a member of the family that the substrate frame now places beside the five numerical projections as the sixth — the one whose honest measure is a relation, because its receiver is never canonical.

7.1 Counterarguments and Responses

Objection 1 (the bridge to model-building is speculative and does not belong in a classification). *Response.* The bridge is marked as a placement, not proved here, and its development is reserved for Part III; what the present article establishes is only that a receiver organized for persistence must reconstruct the differences it tracks, which is the minimal sense of "model." Including the bridge in the classification is appropriate because the receiver axis is where a subsystem first enters the account as the registrant, and a classification that stopped without noting where its last axis leads would misrepresent the architecture of the volume.

Objection 2 (Maxwell's demon belongs to the cost axis, so raising it here is redundant). *Response.* It is raised here only to identify the demon *as a receiver* — a system that acts on registered differences — and to hand its thermodynamic bill to the cost axis, where the resolution lives. The point is the placement: the demon is the receiver axis meeting the cost axis, and naming the junction from the receiver side completes the map of how the six axes touch. The resolution is not attempted here; the junction is.

7.2 Section Result and Implications

We have marked what the receiver axis cannot see — the source of its receiver, and the demon's bill it hands to the cost axis — and shown that its principal border opens onto Part III, since a receiver organized for persistence is a model-building subsystem (Proposition 7.1). We located Bateson, MacKay, Dretske, and Floridi as the authors of the receiver-relative theories. The implication, closing Part II, is that the classification of the projections is complete, and that the volume now turns from the facets of the world to the subsystems that model it.

8 Epistemic Status: Proved, Imported, Affirmed

In the manner of the series, we close with an explicit ledger. *Proved (here)* marks results established in this article; *Imported* marks results from prior articles; *Affirmed (standard)* marks established work the article locates without re-deriving; *Deferred* marks claims handed to a later article or part.

Claim	**Status**	**Locus**
Substrate–projection apparatus; channel reduction	Imported	Kriger [5, 6]
Receiver reduction: register differences that make a difference	Proved (here)	Def. 3.1
No canonical receiver ⇒ no canonical functional	Proved (here)	Prop. 4.1
Channel axis = receiver axis at the maximal observer	Proved (here)	Prop. 5.1
Amount and kind of registered information are receiver-relative	Proved (here)	Fig. 1
Meaning = receiver-relative selection Shannon bracketed	Proved (here)	Prop. 6.1
Receiver organized for persistence is a model-builder	Proved (here; placement)	Prop. 7.1
"Difference that makes a difference" (general form)	Affirmed (standard)	Bateson [1]
Selective information / readiness to act	Affirmed (standard)	MacKay [11]
Information as knowledge-producing (epistemic)	Affirmed (standard)	Dretske [3]
Semantic information (meaningful data)	Affirmed (standard)	Floridi [4]
Shannon's bracketing of semantics	Affirmed (standard)	Shannon [12]
Maxwell's demon as receiver; bill to cost axis	Deferred	Article 9
Model-building subsystems (recover, predict, simulate)	Deferred	Part III

The ledger marks the division of labor. The article *imports* the substrate–projection apparatus and the channel reduction; it *proves* the structural facts about the receiver reduction, the absence of a canonical functional, the channel axis as the maximal-receiver limit, and meaning as receiver-relative selection; it *affirms* the receiver-relative traditions it locates; and it *defers* the demon's resolution to the cost axis and the theory of model-building subsystems to Part III.

9 Discussion

The result of this article is a placement, the sixth and last of the principal projections, and with it the classification of Part II is complete. Bateson's information is located on the receiver axis as the pragmatic projection: the receiver reduction retains the differences that make a difference to a given receiver and discards the rest, and because no receiver is canonical, the axis carries a family of receiver-relative measures rather than a single number. The placement makes two contributions. It shows the channel axis to be the receiver axis at the maximal observer, so that Shannon entropy is the maximal-receiver limit of receiver-relative information and the receiver axis proper is the realm of particular, non-maximal receivers. And it locates meaning — the kind of difference a receiver registers — as the receiver-relative selection that Shannon's maximal observer averages away, dissolving the worry that meaning lies beyond the mathematical theory or beyond physics: meaning is a projection of the substrate along the receiver axis, non-numerical because no receiver is canonical.

The placement completes the architecture of Part II and reveals its shape. Six axes have been located — channel, description, regularity, macrostate, cost, receiver — and five of them carry numbers because their reductions admit a canonical choice once a single parameter is fixed: a source model, a machine, a cut, a coarse-graining, a unit. The sixth carries no number because its parameter, the receiver, has no canonical value, and its information is accordingly a relation. The five numerical axes are the facets a generic registrant can read off the world; the receiver axis is where a particular registrant enters, the system for which some differences matter and others do not. And at its border that registrant becomes a model-builder, which is where the volume now turns. Part II asked what the known measures are measures of, and answered: projections of one differentiated substrate, along axes a choice of reduction selects. Part III asks what a subsystem does with such projections, and the receiver — the last projection and the first model-builder — is the hinge between the two.

9.1 Limitations

Three limitations are material. *First*, the article imports the substrate–projection apparatus and the channel reduction from Articles 3 and 5 and does not re-establish them. *Second*, the figure measures the receiver-relative *amount* with the channel-axis mutual information, which is legitimate once the receiver is fixed but does not, and cannot, capture the *kind* numerically; the kind is displayed by exhibiting distinct receivers, not by a single functional, consistent with the article's central claim that no such functional exists. *Third*, the bridge to model-building subsystems (Proposition 7.1) is stated as a placement and argued only in outline; its development, including the precise sense in which a receiver maintains a model, is the work of Part III and is not discharged here.

9.2 Future Directions

The immediate sequel is Part III, which takes up the model-building subsystem the receiver axis terminates in: the lossiness of projection and the model's recovery, the arrow of time and the incompleteness of measurement as the model-builder meets them, and the relocation of the simulation into the model. Beyond the volume, the family $\{\mu_R\}_R$ of receiver-relative measures invites a precise question continuous with the completeness conjecture of Article 3: whether the space of receivers carries a structure — an ordering by registration, a lattice of coarsenings and feature-selections — whose morphisms could be classified alongside the inter-axis maps, so that the relations among receiver-relative measures become as orderly as the relations among the numerical axes.

That question would turn the receiver's free choice into a position in a classified space, and is left to the work that can address it.

Acknowledgments

The author thanks colleagues of the Information Physics Institute and the Institute of Integrative and Interdisciplinary Research for sustained discussion of the foundations of the series.

References

[1] G. Bateson. *Steps to an Ecology of Mind*. Chandler Publishing, San Francisco, 1972.

[2] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Wiley-Interscience, Hoboken, NJ, 2nd edition, 2006.

[3] F. I. Dretske. *Knowledge and the Flow of Information*. MIT Press, Cambridge, MA, 1981.

[4] L. Floridi. *The Philosophy of Information*. Oxford University Press, Oxford, 2011.

[5] B. Kriger. The projection theorem: how known information measures descend from the substrate. The Information Substrate Theory (IST), Volume VII, Article 3. Zenodo, 2026.

[6] B. Kriger. The axis of the channel: Shannon information as a substrate projection. The Information Substrate Theory (IST), Volume VII, Article 5. Zenodo, 2026.

[7] B. Kriger. The axis of description: Kolmogorov complexity and the algorithmic projection. The Information Substrate Theory (IST), Volume VII, Article 6. Zenodo, 2026.

[8] B. Kriger. The axis of regularity: effective complexity and the structured projection. The Information Substrate Theory (IST), Volume VII, Article 7. Zenodo, 2026.

[9] B. Kriger. The axis of the macrostate: thermodynamic entropy as a substrate projection. The Information Substrate Theory (IST), Volume VII, Article 8. Zenodo, 2026.

[10] R. Landauer. Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5(3):183–191, 1961.

[11] D. M. MacKay. *Information, Mechanism and Meaning*. MIT Press, Cambridge, MA, 1969.

[12] C. E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423, 1948.

A Code for Figure 1

The following Python script generates Figure 1. The source X is uniform on the 64 values (A, B), $A, B \in \{0, \ldots, 7\}$, so $H(X) = 6$ bits. Panel (a) plots $I(X; Q_k(X)) = \log_2 k$ for receivers of resolution k; panel (b) plots $I(X; R(X))$ for receivers organized for different features.

```
import numpy as np
HX = 6.0                                  # H(X) for X uniform on 64 = (A,B)

# (a) amount vs receiver resolution: I(X; Q_k(X)) = log2(k) for equal bins
ks = np.array([2, 4, 8, 16, 32, 64])
I_amount = np.log2(ks)                     # 1,2,3,4,5,6 bits -> channel limit 6

# (b) kind, by what the receiver registers (exact for uniform X):
I_kind = {"shape A":   3.0,               # I(X; A) = H(A) = 3 bits
          "colour B":  3.0,                # I(X; B) = H(B) = 3 bits
          "parity A^B":1.0,                # I(X; (A+B) mod 2) = 1 bit
          "full (A,B)":6.0}                 # I(X; (A,B)) = H(X) = 6 bits
# shape and colour: equal amount (3 bits), different kind.
```

The Atlas Projection: Where the Substrate Axes Materialize across the Scales of the SPCA Programme

An empirical hinge: each abstract reduction axis, located as a concrete object of the fifteen-scale Atlas of the prior volumes

Boris Kriger[1,2]

[1] Information Physics Institute
boris.kriger@informationphysicsinstitute.net

[2] Institute of Integrative and Interdisciplinary Research, Toronto
boriskriger@interdisciplinary-institute.org

ORCID: https://orcid.org/0009-0001-0034-2903

The Information Substrate Theory (IST) — Volume VII, Article 10a
Empirical hinge between Part II (The Projections) and Part III (The Bridge)

Abstract

The classificatory part of this volume has placed six measures of information on six axes of the substrate, each as a reduction with a functional and a characteristic indeterminacy. Stated so, the classification could be read as floating free of physics — an abstract taxonomy with no contact with the concrete results of the programme that precedes it. This article is the hinge that prevents that reading. It does for the whole volume what the prior six volumes did for the cosmos: it shows that each abstract axis materializes, at a definite scale of the fifteen-scale SPCA Atlas, as a concrete physical object already derived and measured. The description axis materializes at the baryon scale as the nucleon's self-consistent topological closure and at the atomic scale as the Banach contraction of an excited spectrum to its stable fixed point. The regularity axis materializes at the biomolecular and cellular scales as the knot topology of DNA and as the deviation Δ from the Banach attractor, read as the volume of control information a living system carries. The channel and macrostate axes materialize at the cosmological scale in the cosmic web, where an infodynamic analysis of DESI DR1 reports a Zipf exponent near 1.29 and a nearest-neighbour mutual information near 0.31 bit. The cost axis materializes in the infodynamics of large-scale processing; the receiver axis in the biosphere as a model-building subsystem; and the global topology $\mathbb{RP}^3$ as the limiting register within which conservation and CPT act as integrity constraints on differentiatedness. Throughout, the article observes the discipline of Articles 2 and 4: it claims that the projections are *realized* at these scales, not that the substrate *is* any of the measures. The result is a single decoder key binding Volume VII to the six volumes it completes.

Keywords: information substrate; SPCA; scale projection; nucleon topology; DNA knots; effective complexity; cosmic web; DESI; $\mathbb{RP}^3$; conservation laws; foundations of physics.

Contents

1 Introduction

A classification can be true and still feel weightless. Part II of this volume placed the known measures of information on six axes of the substrate — channel, description, regularity, macrostate, cost, receiver — and showed each to be a reduction with its own functional and its own characteristic indeterminacy [14, 16–19, 21]. Every step was careful, but a reader closing Article 10 might fairly ask what any of it has to do with the concrete physics of the preceding six volumes: the nucleon computed as a topological soliton, the cosmic web fitted to survey data, the cell analysed by the structural–persistence method. If the axes of Volume VII are merely a tidy way of sorting other people's theories of information, they add a vocabulary and nothing more.

This article answers that the axes are not a vocabulary laid over the physics but a decoder for it. The six prior volumes built an Atlas of structure across fifteen scales, from the sub-nucleonic to the cosmological, using one method — structural–persistence and coherence analysis, with its Banach fixed points, its deviation Δ from the attractor, and its classification of structure into the blindly dissipative, the intermediate, and the coded [5]. The claim here is that each abstract reduction axis of Volume VII *materializes*, at a definite scale of that Atlas, as a concrete object the prior volumes already derived and, where possible, measured. The description axis is the nucleon and the atom; the regularity axis is the knotted DNA molecule and the living cell; the channel and macrostate axes are the cosmic web; the cost axis is the infodynamics of large-scale processing; the receiver axis is the biosphere; and the global topology $\mathbb{RP}^3$ is the register within which the whole stands. Read this way, Volume VII is the key to the Atlas, and the Atlas is the empirical content of Volume VII.

A discipline governs the whole article, and it is the discipline of Articles 2 and 4 [13, 15]. To say that the description axis materializes as the nucleon is to say that the shortest-specification facet of differentiatedness is *realized*, at the baryon scale, in a definite physical object; it is not to say that the nucleon *is* a Kolmogorov program, or that the substrate *is* a code, or that the universe *is* a computer. Those are the reification errors Article 2 diagnosed and Article 4 ruled ill-typed, and this article does not commit them. When an infodynamic analysis of survey data reports a Zipf-like exponent and a nearest-neighbour mutual information, those are empirical regularities of the cosmic web that are *consistent with* the channel and macrostate projections being realized there; they are not a proof that the cosmos is a text. The distinction between "a projection is realized at this scale" and "the substrate is this measure" is the load-bearing line of the article, and it is drawn explicitly at every correspondence. With that line held, the mapping binds Volume VII to its predecessors without overclaiming, and that is the contribution: a single, honest decoder key, not a new metaphysics.

2 Logical Structure of the Argument

The article is a structured map, and its steps are correspondences, not derivations.

Step 1 (Section 3): The scale-projection principle.
We state what it means for an axis to materialize at a scale (Principle 3.1) and, with equal care, what it does not mean — the anti-reification guard inherited from Articles 2 and 4.

Step 2 (Section 4): The description axis materialized.
The nucleon as a self-consistent topological closure (baryon scale) and the atom as a Banach contraction to its stable spectrum (atomic scale).

Step 3 (Section 5): The regularity axis materialized.
The knot topology of DNA (biomolecular scale) and the deviation Δ from the Banach attractor as control information in the living cell (cellular scale).

Step 4 (Section 6): The channel and macrostate axes materialized.
The cosmic web as a coarse-grained macrostate and a channel of neighbouring-cell statistics, with the infodynamic regularities of DESI DR1 read under the anti-reification guard.

Step 5 (Section 7): The cost and receiver axes, and the global register.
The infodynamics of large-scale processing (cost), the biosphere as a model-building subsystem (receiver), and the $\mathbb{RP}^3$ topology as the integrity register of the block.

Step 6 (Section 8): The decoder key.
A single figure and table binding the axes to the Atlas scales, objects, loci, and status.

Step 7 (Section 9): Epistemic status.
A ledger marking each correspondence as established in a prior volume, an empirical anchor reported there, an interpretive correspondence, or a conjecture.

The dependencies are not deductive but cartographic: each section locates one or two axes at their scales, Section 8 assembles the locations into one key, and Section 9 grades them. The thesis is that the abstract classification of Part II is the empirical content of the prior volumes seen through one lens, and that saying so requires neither new physics nor the reification the volume forbids.

3 The Scale-Projection Principle

3.1 What materialization means

We state the principle that licenses the rest of the article.

Principle 3.1 (Scale projection of the Atlas)**.** An axis of the substrate *materializes* at a scale of the Atlas when the object the prior volumes derived at that scale is the physical realization of that axis's reduction: the object is what one obtains when the substrate's differentiatedness is read, at that scale, along that axis. Materialization is realization of a facet at a scale, not identification of the substrate with the facet's measure.

The principle has a positive and a negative content, and both matter. Positively, it asserts a correspondence: the nucleon is what the description facet looks like at the baryon scale, the cosmic web is what the macrostate facet looks like at the cosmological scale, and so on. The correspondence is substantive because the objects were derived independently, in the prior volumes, by physical methods that did not presuppose the axes of Volume VII; that they line up, one object per axis per scale, is the content worth stating. Negatively, the principle refuses the inference from realization to identity. That the description facet is realized as the nucleon does not make the nucleon a program, the substrate a code, or the cosmos a computation; those are precisely the reifications Article 2 diagnosed as confusing a projection with the projected, and Article 4 marked ill-typed [13, 15].

3.2 The anti-reification guard, stated once for all

Because the temptation is strong and recurrent, we state the guard explicitly and apply it silently thereafter. At each scale we will say that a projection is *realized*; we will not say that the substrate *is* the measure. The empirical regularities the prior volumes report — a topological invariant protecting the nucleon, a knot index organizing a macromolecule, a Zipf-like law and a mutual information in the galaxy distribution — are features of the realized objects, consistent with the corresponding projections being realized at those scales. They are evidence that the Atlas and the axes describe the same world; they are not evidence that the world is, at bottom, any one of the measures. The substrate carries no projection value (Article 1); the scales carry the realizations; and the difference between the two is the whole of the discipline this article keeps.

3.3 Section Result and Implications

We have stated the scale-projection principle (Principle 3.1) and the anti-reification guard that bounds it. The implication is that each following section may assert a correspondence between an axis and an Atlas object with a clear conscience, provided it claims realization and not identity. We begin at the smallest scales.

4 The Microscale: the Description Axis Materialized

4.1 The nucleon as a self-consistent closure

The description axis measures the shortest specification of an individual, fixed up to a constant by the choice of machine and characterized by uncomputability and incompressibility [17]. Its physical realization at the baryon scale is the nucleon as a self-consistent topological closure. In the prior volumes the nucleon is computed not as a structureless point but as a chiral soliton whose baryon number is a topological invariant — a winding on the closed configuration space — so that the particle is, in the language of the present volume, an irremovable compressed remainder: a configuration the universe's differentiatedness cannot shed because a conserved integer protects it [1, 4, 6, 7]. On the description axis this is exactly the incompressible individual of Article 6: the nucleon is a "physical bit" in the precise sense that its specification cannot be shortened below the topological minimum, and its persistence is the stability of that minimum under the dissipative dynamics analysed by the structural–persistence method.

The correspondence is sharp and it is bounded. Sharp, because the topological protection of the baryon number is the physical analogue of incompressibility: just as no program shorter than x itself describes a random x, no configuration simpler than the closed soliton carries unit baryon number, and the conserved integer is what forbids the shortcut. Bounded, by the guard of Section 3: the nucleon *realizes* the description facet at the baryon scale; it is not *literally* a Kolmogorov-minimal program, and the soliton's stability is a fact of field theory that the description axis interprets, not replaces.

4.2 The atom as a Banach contraction

One scale up, the description axis is realized again, now dynamically. An excited atom relaxes to its ground configuration by radiating, shedding energy to the vacuum until it reaches a stable member of its spectrum. In the structural–persistence method this relaxation is a contraction mapping: the excited states are carried, step by step, toward a fixed point x^* of the spectrum, the stable orbital, which the dynamics approaches and does not overshoot [1, 5]. On the description axis this is the shedding of algorithmic excess: the atom discards the differentiatedness it cannot stably specify, dropping it into the radiation field, until what remains is the irreducible description of the stable state. The Banach fixed point is the atomic analogue of the incompressible remainder — the configuration that no further contraction shortens.

4.3 Counterarguments and Responses

Objection 1 (calling the nucleon a "bit" is just the reification the article disavows). *Response.* The guard distinguishes the two readings. To say the nucleon realizes the description facet — that the baryon scale is where incompressibility takes physical form as a topologically protected minimum — is licensed by Principle 3.1. To say the nucleon *is* a bit, a program, a piece of code, is the reification Article 2 forbids. The article asserts the first and denies the second; "physical bit" is used as a located metaphor for incompressible persistence, flagged as such, not as an ontological claim.

Objection 2 (the soliton and the contraction are field-theoretic results that

owe nothing to the description axis, so the correspondence is decorative). *Response.* The correspondence is interpretive, not generative, and is presented as such: the prior volumes derived the soliton and the contraction by physical methods, and the description axis names the facet they realize. The value of the naming is the binding it effects — it shows the baryon and atomic scales to be realizations of the same facet that, at other scales, appears as program length — not a claim to have re-derived the physics. A decoder adds no signal; it makes the signal legible across scales, which is its purpose.

4.4 Section Result and Implications

The description axis is realized at the baryon scale as the nucleon's topologically protected closure and at the atomic scale as the Banach contraction to a stable spectrum (Sections 4.1–4.2). The implication is that incompressibility and the irreducible remainder, abstract on the description axis, are concrete and computed at the smallest scales of the Atlas. We ascend to the scales of structure.

5 The Mesoscale: the Regularity Axis Materialized

5.1 Knotted DNA as molecular regularity

The regularity axis measures the structured part of an object — the model in the two-part code, the minimal sufficient statistic — and is non-monotonic, large only for the structured and small for both the ordered and the random [18]. Its physical realization at the biomolecular scale is the knot topology of DNA. In the prior volumes the topological constraints on biomacromolecules — the knot and link types a strand can carry, recognized by the structural–persistence method as invariants of molecular folding — are read as a dimensionless control code: not the random sequence itself, which is the residue, and not a trivial repetition, which would be the ordered extreme, but the organizing topology that coordinates the molecule's dissipative processing [2, 5]. On the regularity axis this is precisely the structured middle of the non-monotonic signature: the knot index is the model, the sequence detail is the residue, and effective complexity is the complexity of the former.

5.2 The living cell and the deviation Δ as control information

At the cellular scale the regularity axis acquires a quantitative handle. The structural–persistence method measures, for a biological network, its deviation Δ from the Banach attractor — the distance by which a living system departs from the blind dissipative flattening that an inert system would undergo [5]. In the present volume's terms, Δ is a measure of the volume of control information the system carries: an inert system relaxes to the attractor (the ordered extreme, little effective complexity), while a living system holds itself off the attractor by the structure encoded in its genome and regulatory network, and the size of that departure is the size of its control program. The more a system deviates from the blind flattening of the dissipative class, the higher the effective complexity of its internal model. On the regularity axis this is the statement that Δ is an empirical proxy for K_{eff} at the cellular scale: a measured distance standing in for the model size the axis defines but cannot, in general, compute.

Remark 5.1 (Why Δ is a proxy, not the measure)**.** Effective complexity is uncomputable (Article 7), so Δ cannot be it exactly; Δ is a computable, measured distance that tracks the structured share of a living system's organization, in the way that a compression ratio tracks Kolmogorov complexity from above. The correspondence is that Δ realizes, at the cellular scale, the quantity the regularity axis names — the volume of structure that holds a system off the dissipative attractor — not that Δ equals K_{eff}. The guard of Section 3 applies: the cell realizes the regularity facet; it is not a literal minimal sufficient statistic.

5.3 Counterarguments and Responses

Objection 1 (identifying Δ with control information smuggles in teleology). *Response.* It does not: Δ is a measured distance from a dynamical attractor, defined without reference to purpose, and "control information" names the structural role that distance plays — holding the system off the flattening an inert system would suffer — not a goal the system pursues. The receiver axis (Article 10) is where the language of function and persistence properly enters; here Δ is a regularity-axis quantity, the structured share, and its interpretation as control is structural, not teleological.

Objection 2 (knot topology is a property of the molecule, not "information"). *Response.* It is a property of the molecule that realizes the regularity facet: a dimensionless invariant that organizes the molecule's behaviour and is neither its random sequence nor a trivial pattern. That is exactly what the regularity axis measures — structure as distinct from residue and from order — and naming the knot index a control code is the located reading Principle 3.1 licenses, bounded by the guard against calling the molecule a program.

5.4 Section Result and Implications

The regularity axis is realized at the biomolecular scale as the knot topology of DNA and at the cellular scale as the deviation Δ, a measured proxy for the control information that holds a living system off the dissipative attractor (Sections 5.1–5.2, Remark 5.1). The implication is that the non-monotonic effective complexity of Article 7 is, at the scales of life, a measured quantity of the Atlas. We ascend to the largest scales.

6 The Macroscale: the Channel and Macrostate Axes Materialized

6.1 The cosmic web as a coarse-grained macrostate

The macrostate axis measures the coarse-grained entropy of a system, $S = k_{\mathrm{B}}H$ of the microstate distribution, with the second law as the lossiness of the coarse-graining [19]. Its physical realization at the cosmological scale is the cosmic web — the distribution of galaxies, voids, and filaments mapped by surveys such as DESI and BOSS. In the prior volumes this distribution is not a random outcome of gravitational instability but a structured fixed point of the universe's large-scale dynamics under the vacuum coupling, the self-organizing foam the structural–persistence method recognizes at the top of the Atlas [3, 8]. On the macrostate axis the web is a coarse-grained macrostate: a binning of the matter field whose entropy is the registered differentiatedness of the cosmological coarse-graining, and whose growth toward equilibrium is the cosmological face of the second law treated in Article 8.

6.2 The channel axis and the infodynamic regularities of DESI DR1

The same distribution realizes the channel axis, and here the prior programme reports a measurement worth stating with care. An infodynamic analysis of the DESI DR1 field, treating the binned galaxy distribution as a stream of symbols, reports a Zipf-like frequency law with exponent near $\beta \approx 1.29$ and a nearest-neighbour mutual information near 0.31 bit between adjacent cells [10, 20]. On the channel axis the mutual information is exactly the registered dependence between neighbouring regions — the channel-statistics of the web, the amount by which one cell's state reduces uncertainty about its neighbour's — and the Zipf-like law is a regularity of the frequency distribution of cell states.

Here the guard of Section 3 does its most important work. These regularities are *consistent with* the channel and macrostate projections being realized in the cosmic web: a structured macrostate with non-trivial neighbour dependence is what those axes, realized at the cosmological scale, would display. They are *not* a demonstration that the universe is a text, a language, or a computation. A Zipf-like exponent and a positive mutual information are statistical features shared by many structured fields, linguistic and non-linguistic alike; reading them as proof that the cosmos *is* code would be the reification Article 2 diagnosed and Article 4 forbade. What the measurement supports is the weaker and secure claim that the cosmic web carries the channel and macrostate facets in measurable form — that the abstract axes are realized, at the largest scale, in numbers a survey can return. The hypothesis that the universe is a computational structure, raised philosophically in Article 4, is neither proved nor needed here; the empirical content is the realization of the projections, not the identity of the substrate with a computer.

6.3 Counterarguments and Responses

Objection 1 (a Zipf law and a mutual information in DESI prove the universe is a language). *Response.* They do not, and the article is explicit that they do not. Zipf-like laws and positive neighbour mutual information arise in a wide range of structured systems without language, and the inference from such regularities to "the cosmos is a text" is exactly the projection-for-substrate confusion the volume forbids. The regularities are reported as empirical features of the web consistent with the channel and macrostate projections being realized there; the linguistic vocabulary, if used at all, is metaphor under the guard, not a claim.

Objection 2 (if the regularities prove nothing about the substrate, why report them). *Response.* Because they are evidence that the Atlas and the axes describe one world: that the cosmological scale, derived in the prior volumes by gravitational and vacuum dynamics, returns precisely the channel-statistics and macrostate structure that the abstract axes would predict it to carry. That the two independent descriptions meet in a measured number is the binding the article exists to make. The numbers do real work — they anchor the macroscale of the Atlas to the channel and macrostate axes — without doing the illegitimate work of reifying the substrate.

6.4 Section Result and Implications

The macrostate and channel axes are realized at the cosmological scale in the cosmic web: a coarse-grained macrostate whose neighbour statistics, reported by the infodynamic analysis of DESI DR1 as a Zipf-like law and a 0.31-bit mutual information, are consistent with the projections being realized there, without reifying the substrate (Sections 6.1–6.2). The implication is that the two most physical of the numerical axes are anchored at the top of the Atlas. We complete the map with the remaining axes and the global register.

7 The Cost and Receiver Axes, and the Global Register

7.1 The cost axis: infodynamics of large-scale processing

The cost axis prices information physically, through Landauer's bound $k_{\mathrm{B}}T \ln 2$ per erased bit, and is the subject of the volume's infodynamics article (Article 9). Its realization in the Atlas is the energetics of large-scale information processing: the analysis of entropy flow and the Landauer cost of algorithmic compression in physical systems, which the prior programme develops both for engineered systems and, by extension, for the cosmological processing the second law tracks [10, 20]. On the cost axis these are the thermodynamic prices of the differentiations the other axes count, and the macrostate axis met this junction already at Maxwell's demon (Article 8): the demon is a receiver that pays the cost axis for the differences it registers. The cost axis thus binds the macrostate and receiver axes through the energetics the infodynamics article quantifies.

7.2 The receiver axis: the biosphere as a model-building subsystem

The receiver axis carries no canonical number, only a family of receiver-relative measures, and its border opens onto the model-building subsystem (Article 10). Its realization in the Atlas is the biosphere. The prior programme estimates the biosphere's relative contribution to the effective complexity of the universe — the share of the cosmos's structured information that life and, latterly, synthetic cognition carry — and finds it a small but distinguished contribution, concentrated in systems that hold themselves off the dissipative attractor [11, 12]. On the receiver axis the biosphere is the realized receiver at planetary scale: the collection of subsystems organized to register the differences that bear on their persistence, and so to model their world. This is the realization that hands Part II to Part III, for the biosphere is the Atlas's instance of the model-building subsystem the receiver axis terminates in [21].

7.3 The global register: $\mathbb{RP}^3$ topology and the integrity constraints

Beneath every scale of the Atlas lies one global fact: the closure of the cosmos on a compact manifold, $\mathbb{RP}^3 \times \mathbb{R}$ in the prior programme's topology [3, 9]. On the reading of the present volume this global topology is the limiting register within which all the projections are realized — the closed configuration space that makes the substrate's differentiatedness a bounded, conserved totality rather than an open and leaking one. Within such a register the conservation laws and CPT invariance acquire a structural role: they are the constraints that prevent any quantum of differentiatedness from vanishing without trace from the block, the rules by which the totality keeps its books. The analogy with the integrity checks of a data register — "checksums" that no bit be silently lost — is a heuristic, used in quotation marks and not as a claim that the cosmos is a memory device; the substantive content is that closure plus conservation makes the block a conserved totality of differentiatedness, which is the global precondition the projections presuppose.

7.4 Counterarguments and Responses

Objection 1 (the biosphere's "contribution to the effective complexity of the universe" is not measurable and the estimate is speculative). *Response.* The estimate is presented in the prior volumes as an order-of-magnitude assessment under stated assumptions, not a precise measurement, and it is cited here at that status (Section 9). What the correspondence requires is only that the biosphere realize the receiver facet — that it be the planetary collection of model-building subsystems — which is secure independently of the numerical estimate; the estimate quantifies the realization without being the basis of the correspondence.

Objection 2 (calling conservation laws "checksums" is the computational reification again). *Response.* It is flagged as a heuristic in quotation marks for exactly that reason. The claim made without metaphor is structural and modest: closure of the cosmos on a compact manifold makes differentiatedness a bounded, conserved totality, and conservation and CPT are the constraints that keep it so. The checksum image conveys the role — nothing lost without trace — to a reader, but the article does not assert that the universe performs error-correction or is a register in any literal sense. The guard holds here as everywhere.

7.5 Section Result and Implications

The cost axis is realized in the infodynamics of large-scale processing, the receiver axis in the biosphere as a model-building subsystem, and the global $\mathbb{RP}^3$ topology as the integrity register within which all projections stand (Sections 7.1–7.3). The implication is that the map is complete: six axes and one global register, each located in the Atlas. We assemble the key.

8 The Decoder Key

Figure 1 places the realizations on a single scale ladder, and the table that follows records, for each axis, its Atlas object, the scale, the prior-volume locus, and the status of the correspondence.

Axis	Atlas object	Scale	Prior locus
Description K	Nucleon: self-consistent topological closure	baryon ($\sim$fm)	Vol. I, Vol. V [1, 4, 6, 7]
Description K	Atom: Banach contraction to stable spectrum x^*	atomic ($\sim$Å)	Vol. I, Vol. VI [1, 5]

continued on next page

Decoder key (continued)

Axis	Atlas object	Scale	Prior locus
Regularity K_{eff}	Knotted DNA: molecular knot topology as control code	biomolecular	Vol. VI (SPCA) [2, 5]
Regularity K_{eff}	Living cell: deviation Δ from attractor = control information	cellular	Vol. VI (SPCA) [5]
Macrostate S	Cosmic web: coarse-grained matter macrostate	cosmological	Vol. III [3, 8]
Channel H	Cosmic web: neighbour mutual information ($\approx$ 0.31 bit), Zipf $\beta \approx 1.29$	cosmological	Article 9 [10, 20]
Cost (Landauer)	Infodynamics of large-scale processing	cross-scale	Article 9 [10, 20]
Receiver (Bateson)	Biosphere: model-building subsystem	biospheric	[11, 12]
Global register	$\mathbb{RP}^3$ closure; conservation, CPT as integrity constraints	global	Vol. III [3, 9]

The key is the article in one view: the abstract axes down the left, the concrete Atlas objects across the middle, the scales and prior loci on the right. Read across any row and the same fact appears — a facet of differentiatedness, abstract in Part II, realized at a scale in the prior volumes. Read down the table and the Atlas reassembles, from the nucleon to the cosmic web, as the empirical content of the classification.

9 Epistemic Status: Established, Anchored, Interpretive, Conjectured

This article asserts correspondences, and honesty requires grading them. *Established (prior volume)* marks an object derived and defended in a prior volume; *Empirical anchor (reported)* marks a measured or estimated number reported there; *Interpretive correspondence* marks the reading of that object as the realization of an axis, which is the present article's contribution and carries the anti-reification guard; *Conjectured* marks a correspondence offered as plausible but not secured.

Claim	Status	Locus
Nucleon as topologically protected soliton (baryon number invariant)	Established (prior volume)	Kriger, Vol. V [4], Kriger [6]

continued on next page

Status ledger (continued)

Claim	Status	Locus
Atomic relaxation as Banach contraction to x^*	Established (prior volume)	Kriger, Vol. VI [5]
DNA knot topology; deviation Δ from attractor	Established (prior volume)	Kriger, Vol. VI [5]
Cosmic web as structured fixed point of large-scale dynamics	Established (prior volume)	Kriger, Vol. III [3], Kriger [8]
$\mathbb{RP}^3$ closure of the cosmos	Established (prior volume)	Kriger, Vol. III [3], Kriger [9]
DESI DR1 Zipf $\beta \approx 1.29$; neighbour MI ≈ 0.31 bit	Empirical anchor (reported)	Kriger [10], Kriger, IST 9 [20]
Biosphere contribution to cosmic effective complexity	Empirical anchor (estimate)	Kriger [11, 12]
Nucleon/atom realize the description axis	Interpretive correspondence	§4
DNA/cell realize the regularity axis; Δ proxies K_{eff}	Interpretive correspondence	§5
Cosmic web realizes the channel and macrostate axes	Interpretive correspondence	§6
Biosphere realizes the receiver axis	Interpretive correspondence	§7.2
Conservation/CPT as "integrity constraints" (heuristic)	Interpretive (heuristic)	§7.3
Substrate is *not* identified with any measure	Guard (held throughout)	Kriger, IST 2 [13], Kriger, IST 4 [15]
Single fixed point binding all scales into one computation	Conjectured	§10

The ledger keeps the article honest. The Atlas objects and their numbers are *established* or *anchored* in the prior volumes; the readings of those objects as realizations of the axes are *interpretive*, the present article's contribution, and they carry the guard against reification at every step; the checksum image is flagged *heuristic*; and the strongest unifying claim — that all scales are positions in one global computation — is marked *conjectured*, not asserted. The one thing the article refuses throughout is the identification of the substrate with any of the measures, which is the guard of Articles 2 and 4.

10 Discussion

The result of this article is a binding. The six axes of Part II, which could be read as an abstract taxonomy, are shown to be the empirical content of the six prior volumes

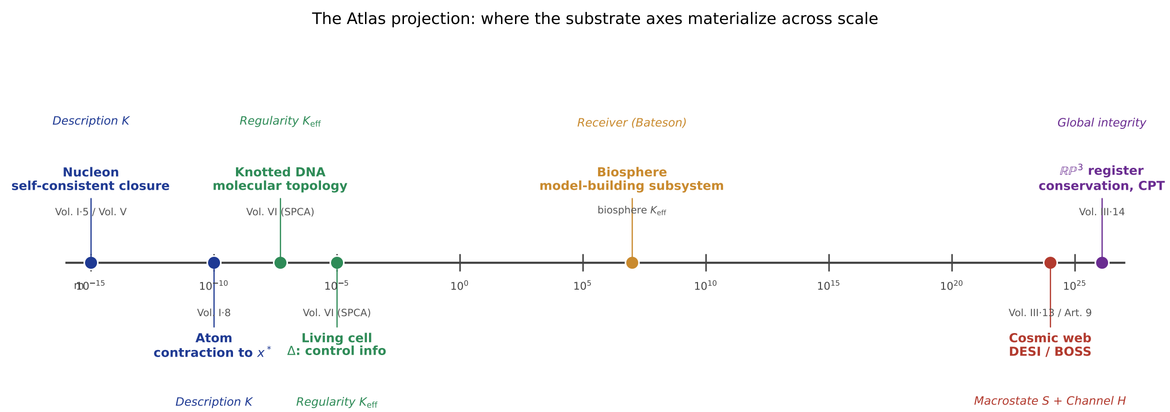

Figure 1: The decoder key as a scale ladder. Each reduction axis of Volume VII is located at the scale of the Atlas where it materializes as a concrete object derived in the prior volumes, spanning roughly forty orders of magnitude of length. The description axis (blue) at the baryon and atomic scales; the regularity axis (green) at the biomolecular and cellular scales; the receiver axis (amber) at the biospheric scale; the channel and macrostate axes (red) at the cosmological scale; and the global $\mathbb{RP}^3$ register (purple) closing the block. The figure is a map of where projections are realized, not a claim that the substrate is any of the measures.

seen through one lens: each axis materializes, at a definite scale of the fifteen-scale Atlas, as a concrete object the prior volumes derived and, where possible, measured. The description axis is the nucleon and the atom; the regularity axis is the knotted molecule and the living cell; the channel and macrostate axes are the cosmic web; the cost axis is the infodynamics of processing; the receiver axis is the biosphere; and the global $\mathbb{RP}^3$ topology is the register within which all of them stand. The decoder key of Section 8 is the whole of it in one view, and it does the work the volume needed: it shows that Volume VII does not lead away from physics into philosophy but back into the physics of the prior volumes, furnishing a single information-theoretic key to the scales they mapped.

The binding is made under a discipline, and the discipline is the point as much as the binding. At every correspondence the article claims that a projection is *realized* at a scale and refuses the inference that the substrate *is* the measure. The nucleon realizes incompressibility without being a program; the cosmic web realizes the channel and macrostate facets, and its Zipf-like law and mutual information are consistent with that realization, without making the cosmos a text; conservation and CPT keep the block's books, and the checksum image conveys the role without asserting a register. This is the discipline of Articles 2 and 4, carried into the empirical hinge, and it is what lets the article bind the seven volumes into one programme without inflating a classification into a metaphysics. The seven-volume whole is monolithic not because every scale *is* a computation, but because every scale *realizes*, in its own physics, a facet of the one differentiated substrate — and that is a claim the volume can make and defend.

10.1 Limitations

Three limitations are material. *First*, the correspondences are interpretive: the prior volumes derived the objects by physical methods independent of the axes, and the present article reads them as realizations, a reading that adds legibility across scales but no new physical result, and that a reader may accept the physics while declining. *Second*, the empirical anchors — the DESI infodynamic regularities and the biosphere estimate — are reported from the prior programme at the status given in the ledger, and the article neither re-derives nor independently validates them; their weight is the weight of their sources. *Third*, the strongest unifying picture, that the fifteen scales are positions in one global computation with a single fixed point, is a conjecture the article declines to assert, holding instead to the modest and defensible claim that each scale realizes a facet of one substrate.

10.2 Future Directions

The immediate sequel is Part III, which the receiver axis and the biosphere have already opened: the theory of the model-building subsystem, the lossiness of projection and the model's recovery, and the relocation of the simulation into the model. The decoder key invites one precise question continuous with the completeness conjecture of Article 3: whether the map from axes to scales is itself structured — whether the order of the scales mirrors the order of the inter-axis morphisms, so that ascending the Atlas is traversing the morphism diagram — which would turn the decoder key from a table into a theorem. That question, and the global-computation conjecture the ledger marks, are left to the work that can address them; the present article's task was the binding, and the binding is made.

Acknowledgments

The author thanks colleagues of the Information Physics Institute and the Institute of Integrative and Interdisciplinary Research for sustained discussion of the foundations of the series, and acknowledges the six prior volumes of the programme whose results this article binds to the present classification.

References

[1] B. Kriger. *Volume I: The Local Gravitation of Quantum Vacuum — A Unified Solution to the Dark Sector (αLGQV Theory Monograph)*. IIIR Cosmology and Theoretical Physics, Toronto, 2026. doi:10.5281/zenodo.19027460.

[2] B. Kriger. *Volume II: The Consistent Universe — A Dark-Sector-Free Cosmology*. IIIR Cosmology and Theoretical Physics, Toronto, 2026. doi:10.13140/RG.2.2.29913.28002.

[3] B. Kriger. *Volume III: Predictions Confirmed — NANOGrav, P-ACT, DESI, Big Ring*. IIIR Cosmology and Theoretical Physics, Toronto, 2026. doi:10.13140/RG.2.2.32586.32962.

[4] B. Kriger. *Volume V: Topology of Quantum Chromodynamics and Isotopology*. IIIR Cosmology and Theoretical Physics, Toronto, 2026. doi:10.5281/zenodo.20370383.

[5] B. Kriger. *Volume VI: Structural Persistence and Coherence Analysis — A Topological–Dissipative Method Across Physical Scales*. IIIR Cosmology and Theoretical Physics, Toronto, 2026. doi:10.5281/zenodo.20455316.

[6] B. Kriger. On the topological origin of QCD confinement: why closed color configurations are internally stable. IIIR. ResearchGate, 2026. doi:10.13140/RG.2.2.13146.30404.

[7] B. Kriger. Mass as topology: nuclear configurations as rendered vacuum activity in S^3-closed form. IIIR. ResearchGate, 2026. doi:10.13140/RG.2.2.17655.18085.

[8] B. Kriger. The cosmic web as a self-organising foam: structural laws, local topology, and the limits of global tests on DESI DR1 data. IIIR Cosmology and Theoretical Physics, 2026. doi:10.13140/RG.2.2.26354.67526.

[9] B. Kriger. Dynamic topology of the universe: $\mathbb{RP}^3$ with double counter-rotation as a complete physical model. IIIR Cosmology and Theoretical Physics, 2026. doi:10.13140/RG.2.2.20020.62086.

[10] B. Kriger. Local entropy inversion in large-scale AI systems: Landauer bounds on algorithmic compression. *IPI Letters*, 4(2):13–20, 2026. doi:10.59973/ipil.335.

[11] B. Kriger. Significance of biosphere relative contribution to effective complexity of the universe. IIIR. Zenodo, 2026. doi:10.5281/zenodo.18256772.

[12] B. Kriger. Estimation of the contribution of biospheric and synthetic cognitive systems to the total effective complexity of the universe. IIIR. Zenodo, 2026. doi:10.5281/zenodo.18261863.

[13] B. Kriger. Substrate and projection: why the universe needs no information. IST, Volume VII, Article 2. Zenodo, 2026.

[14] B. Kriger. The projection theorem: how known information measures descend from the substrate. IST, Volume VII, Article 3. Zenodo, 2026.

[15] B. Kriger. The categorial boundary: on predicates that do not cross the block. IST, Volume VII, Article 4. Zenodo, 2026.

[16] B. Kriger. The axis of the channel: Shannon information as a substrate projection. IST, Volume VII, Article 5. Zenodo, 2026.

[17] B. Kriger. The axis of description: Kolmogorov complexity and the algorithmic

projection. IST, Volume VII, Article 6. Zenodo, 2026.

[18] B. Kriger. The axis of regularity: effective complexity and the structured projection. IST, Volume VII, Article 7. Zenodo, 2026.

[19] B. Kriger. The axis of the macrostate: thermodynamic entropy as a substrate projection. IST, Volume VII, Article 8. Zenodo, 2026.

[20] B. Kriger. The axis of cost: Landauer's principle and the infodynamics of DESI DR1. IST, Volume VII, Article 9. Zenodo, 2026.

[21] B. Kriger. The axis of the receiver: Bateson's difference and the pragmatic projection. IST, Volume VII, Article 10. Zenodo, 2026.

Projection Loses, the Model Recovers: The Central Asymmetry and the Opening of the Bridge

The three losses of Part II as one, and how a subsystem recovers structure by inference rather than by inversion

Boris Kriger[1,2]

[1] Information Physics Institute
boris.kriger@informationphysicsinstitute.net

[2] Institute of Integrative and Interdisciplinary Research, Toronto
boriskriger@interdisciplinary-institute.org

ORCID: https://orcid.org/0009-0001-0034-2903

The Information Substrate Theory (IST) — Volume VII, Article 11
Part III: The Bridge — opening article

Abstract

Part II left three observations standing, one on each of three axes. On the channel axis, the data-processing inequality: no operation on a projection recovers information about the source that the projection discarded. On the description axis, the uncomputability of inversion: the shortest preimage of an output is not effectively findable. On the macrostate axis, the second law: under reversible micro-dynamics the coarse-grained entropy rises, and the coarse-graining cannot be undone. This article, opening the bridge part of the volume, shows that these are one fact wearing three faces — physical projection is many-to-one and therefore lossy and irreversible — and then asks how, given that the world cannot unproject, any subsystem comes to know the structure projection discarded. The answer is the central asymmetry of the volume: projection loses, and the model recovers, but the model does not recover by inverting the projection, which is blocked. It recovers by inference. A subsystem maintains an internal surrogate of the preimage, conditioned on the projections it receives and updated defeasibly against the projections still to come; this surrogate recovers the projectable structure progressively, without ever physically recreating a discarded bit. We make the asymmetry precise, show that it violates no irreversibility — because the model is a separate, costly, fallible physical system, not the resurrected source — and demonstrate with a Bayesian subsystem that inference recovers the projectable structure of a hidden state while the discarded fiber remains an irreducible floor. The floor is the residue that no inference removes, and it is the seed of the incompleteness the later bridge articles develop. With the asymmetry in hand, the part's programme is set: the arrow of time, the relocation of the simulation into the model, and the incompleteness of measurement are the model-builder's situation, examined in turn.

Keywords: lossy projection; data-processing inequality; irreversibility; Bayesian inference; model-building subsystem; predictive processing; incompleteness; foundations of physics.

Contents

1 Introduction

The classificatory part of this volume closed with a pattern and a debt. The pattern was that six measures of information are projections of one differentiated substrate along six axes, each a reduction with a functional and a characteristic indeterminacy [6–11]. The debt was three observations the part raised and deferred, one on each of three axes, all of the same shape. On the channel axis, the data-processing inequality says that no processing of a received signal can increase its mutual information with the source: what the channel discarded is gone [7]. On the description axis, the uncomputability of inversion says that the shortest program for an output cannot be found effectively: the minimal preimage is not recoverable by computation [8]. On the macrostate axis, the second law says that under reversible micro-dynamics the coarse-grained entropy rises and the coarse-graining cannot be undone, so the fine structure the macrostate omits is, for the macrostate, lost [10]. Three axes, three losses, deferred to the bridge.

This article opens the bridge by paying the debt in two moves. The first move is unification: the three losses are one fact. A projection is a many-to-one map — the reduction identifies configurations the axis does not distinguish — and a many-to-one map has no inverse as a function. The data-processing inequality, the uncomputability of inversion, and the second law are the channel-, description-, and macrostate-axis expressions of the single circumstance that projection discards a fiber, the set of preimages collapsed to one image, and the fiber does not come back. Physical projection loses, and it loses irreversibly, because the world offers no free process that would un-collapse the fiber.

The second move is the central asymmetry of the volume, and it answers the question the losses provoke. If projection is irreversibly lossy, how does any subsystem — an organism, an instrument, a science — come to know the structure projection discarded? It cannot be by inversion: the world does not run the projection backward, and a subsystem inside the world has no power the world lacks. The answer is that the model recovers what projection lost, but not by inverting it. A model-building subsystem maintains an internal surrogate of the preimage, conditioned on the projections it receives and corrected against the projections still to come; the surrogate is a hypothesis, defeasible and costly to maintain, tested against the stream of further projections. This recovers the projectable structure progressively — not by recreating a discarded bit, which remains discarded, but by inferring the preimage well enough to predict the next projection. Recovery lives in the subsystem, as inference; loss lives in the world, as irreversibility; and the two never meet as inverse operations.

The asymmetry dissolves an apparent paradox and seeds a real limit. It dissolves the paradox of how knowledge is possible under lossy projection: knowledge is the model's inference, not the world's inversion, and inference can constrain the preimage across many projections without ever undoing one. And it seeds the limit the later bridge articles develop: the model's inference, however prolonged, leaves a residue — the part of the fiber no projection ever distinguishes — and that residue is an irreducible floor, the seed of the incompleteness of measurement. We demonstrate the recovery and the floor with a Bayesian subsystem, ground the model-builder in the prior work that shows bounded systems must model [3–5], and set the programme of Part III.

2 Logical Structure of the Argument

The article moves from the unification of the losses to the asymmetry that answers them.

Step 1 (Section 3): The three losses are one.
A projection is many-to-one; the data-processing inequality, the uncomputability of inversion, and the second law are the three axis-faces of the discarded fiber (Proposition 3.1).

Step 2 (Section 4): Physical inversion is blocked.
No physical process recovers the fiber for free; the world does not un-project (Proposition 4.1).

Step 3 (Section 5): The model recovers by inference.
A subsystem maintains a defeasible internal surrogate of the preimage, conditioned on received projections, recovering the projectable structure without inverting the projection (Proposition 5.1, Figure 1).

Step 4 (Section 6): The asymmetry without paradox.
Modeling is not inverse projection; it violates no irreversibility, since the model is a separate, costly, fallible system, and it leaves an irreducible floor (Proposition 6.1).

Step 5 (Section 7): What the bridge opens.
The arrow of time, the relocation of the simulation, and the incompleteness of measurement are the model-builder's situation, set as the programme of Part III.

Step 6 (Section 8): Epistemic status.
A ledger separates what is proved here, what is imported, and what is affirmed or deferred.

The dependencies are linear: Step 2 sharpens Step 1 into irreversibility; Step 3 introduces the subsystem that Step 2 shows cannot invert but can infer; Step 4 makes the asymmetry precise and marks the floor; Step 5 hands the floor and the asymmetry to the rest of Part III. The thesis is that projection loses irreversibly and the model recovers by inference, two operations of different kinds in different places, whose distinctness is the engine of the bridge.

3 The Three Losses Are One

3.1 Projection is many-to-one

The unification rests on a single structural fact about reductions, already implicit in every axis of Part II.

Proposition 3.1 (The discarded fiber, and its three faces)**.** *Each axis-reduction ρ is a many-to-one map: it identifies substrate configurations the axis does not distinguish, sending a set of preimages — the* fiber *— to a single projected value. A many-to-one map has no functional inverse. The three losses of Part II are the axis-specific expressions of this one fact:*

(i) (Channel.) *The data-processing inequality: for a Markov chain $X \to Y \to Z$, $I(X;Z) \le I(X;Y)$, so no processing of the projection Y recovers information about the source X that ρ discarded [7].*

(ii) (Description.) *The uncomputability of inversion: the shortest preimage of an output is upper semi-computable but not computable; the minimal member of the fiber is not effectively findable [8].*

(iii) (Macrostate.) *The second law: the coarse-graining map is many-to-one on microstates, and the coarse-grained entropy — the log-size of the fiber — does not spontaneously decrease, so the fiber is not undone [10].*

Proof. That each reduction is many-to-one is its defining property: the channel reduction identifies outcomes by their statistics, the description reduction by nothing but is read through a many-to-one machine semantics on programs, the macrostate reduction by the coarse-graining partition, each collapsing a fiber to an image (Part II). A many-to-one map has no inverse function, since the image does not determine the preimage. The three named results are then immediate readings of the collapse: *(i)* the data-processing inequality is the statement that post-processing cannot enlarge the preimage information [1]; *(ii)* uncomputability of the minimal preimage is the description-axis inversion obstruction [8]; *(iii)* the non-decrease of fiber-size is the second law [10]. Each is the same fiber, seen along one axis. □

The proposition is the hinge of the bridge. What looked, in Part II, like three separate difficulties — an information-theoretic inequality, a computability obstruction, a thermodynamic law — is one fact about reductions: they collapse fibers, and fibers do not uncollapse. The loss is not a defect of any particular measure but the price of projecting an axis-free substrate onto a single coordinate, and it is the same price on every axis.

3.2 Counterarguments and Responses

Objection 1 (the three results have different proofs and domains, so calling them one fact overreaches). *Response.* They have different proofs because they live on different axes with different apparatus; the claim is not that the proofs coincide but that the proofs establish, each in its idiom, the non-invertibility of a many-to-one reduction. The data-processing inequality bounds information across a channel, the uncomputability bounds effective inversion of a description, the second law bounds the undoing of a coarse-graining; all three are the fiber refusing to uncollapse. Shared content, distinct proofs, is the normal situation of a unifying observation, not an overreach.

Objection 2 (some projections are invertible — a bijection loses nothing — so the unification fails). *Response.* A bijection is the degenerate reduction that distinguishes everything, with singleton fibers and no loss; it is the maximal-observer limit (Article 10) at which the receiver axis collapses to the channel axis and the projection is no projection at all. The unification concerns proper reductions, those that genuinely identify — the ones the axes of Part II are about — and for those the fiber is non-trivial and the loss real. That the trivial reduction loses nothing is consistent with, and indeed sharpens, the claim.

3.3 Section Result and Implications

The three losses of Part II are one fact: projection is many-to-one, collapsing a fiber that has no functional inverse, and the data-processing inequality, the uncomputability of inversion, and the second law are its three axis-faces (Proposition 3.1). The implication is that loss is generic to projection, not special to any axis. We sharpen the loss into irreversibility.

4 Physical Inversion Is Blocked

4.1 The world does not un-project

A many-to-one map has no inverse *function*, but one might hope that some physical process could nonetheless reconstruct the fiber — that the world could run the projection backward. It cannot, and the same three axes say why.

Proposition 4.1 (Physical inversion is blocked)**.** *No physical process recovers the discarded fiber without cost or in general at all:*

(i) (Channel.) *By the data-processing inequality, any physical post-processing of a projection is a further channel, which cannot increase mutual information with the source [1, 7].*

(ii) (Cost.) *Restoring a discarded distinction is a form of un-erasure; by Landauer's principle the erasure that discarded it dissipated* $k_{\mathrm{B}}T \ln 2$ *per bit to the environment, and that environment's record is itself coarse-grained away, so the reconstruction has no free physical source [10, 12].*

(iii) (Macrostate.) *By the second law, the coarse-grained entropy of an isolated system does not spontaneously decrease, so the fiber's collapse is not spontaneously reversed [10].*

Hence the world offers no free, general process that un-projects: physical inversion is blocked.

Proof. Each clause is a direct application. *(i)* A physical operation on the projection is a stochastic map applied after ρ, hence a channel composed with the projection, and the data-processing inequality bounds the composed information by the projection's [1]. *(ii)* The discarding of a distinction is logically irreversible and so, by Landauer, thermodynamically costly; the heat that carried the distinction away disperses into a coarse-grained environment from which it is not freely retrievable [12]. *(iii)* The non-decrease of coarse-grained entropy is the second law (Article 8). Together they leave no free channel, no free energy source, and no spontaneous reversal by which the fiber returns. □

The proposition is the negative half of the asymmetry. It is tempting, faced with a lossy projection, to imagine that a clever enough physical apparatus could undo it — could, given the image, physically produce the preimage. The three axes forbid it in concert: information-theoretically the apparatus is just another channel, bounded by what it received; thermodynamically the lost distinction was paid out as heat into an unrecoverable environment; and dynamically the coarse-graining does not spontaneously reverse. Whatever recovers the fiber, it is not the world running backward.

4.2 Counterarguments and Responses

Objection 1 (microscopic dynamics is reversible, so in principle the fiber can be recovered by reversing the trajectories). *Response.* Microscopic reversibility is exactly what Article 8 invoked, and it does not help: reversing every microscopic trajectory would require knowing the microstate, which is the fiber one is trying to recover, and the coarse-grained access a subsystem actually has does not contain it. In-principle microscopic reversal presupposes the very information the projection discarded; it is not a process available to any subsystem that receives only the projection. The block is on inversion *from the projection*, which is the only thing a subsystem has.

Objection 2 (error-correcting codes recover discarded information routinely, so inversion is not blocked). *Response.* Error correction recovers what redundancy *preserved*, not what projection discarded: a code adds structure to the source before transmission so that the channel's loss falls on the redundancy, leaving the message recoverable. That is preservation against an anticipated loss, not inversion of a completed projection. The fiber a code recovers was never in the fiber that was discarded; it was protected in advance. Where no redundancy was added — the generic projection of Part II — there is nothing for a code to exploit, and the block stands.

4.3 Section Result and Implications

Physical inversion is blocked: no free, general process un-projects, by the data-processing inequality, Landauer's principle, and the second law in concert (Proposition 4.1). The implication is that whatever recovers the discarded structure is not the world inverting the projection. We introduce the thing that does.

5 The Model Recovers by Inference

5.1 The subsystem's surrogate

If the world cannot un-project, recovery must come from elsewhere, and the elsewhere is the model-building subsystem the receiver axis terminated in (Article 10). It recovers not by inversion but by inference.

Proposition 5.1 (Model recovery is inference, not inversion)**.** *A model-building subsystem maintains an internal surrogate M of the preimage — a hypothesis about the substrate structure behind the projections it receives — and updates M on each projection by conditioning, $M_k \propto M_{k-1} \cdot \Pr(y_k \mid \cdot)$, where y_k is the k-th projection. The surrogate recovers the projectable structure progressively: as projections accumulate, M concentrates on the preimages consistent with all of them. This recovery is inference, not inversion: M is a hypothesis maintained in the subsystem and tested against the projections still to come, not a physical reconstruction of any discarded bit. No clause of Proposition 4.1 is violated, because no discarded distinction is physically recreated — only posited and corrected.*

Argument. The update is Bayesian conditioning, which requires only the received projections and a model of how the substrate structure would generate them; it adds no physical channel to the source and so does not contradict the data-processing inequality, which bounds information *about the source carried by the signal*, not the subsystem's posterior *given its model*. The posterior concentrates on the set of preimages consistent with the projections received — the projectable structure — by standard consistency of Bayesian updating under an identifiable model [2]. That the surrogate is a hypothesis and not a reconstruction is what keeps it inside the subsystem: it is realized in the subsystem's own states, costs energy to maintain and update [12], and is fallible, being only as good as its model and its projections. The recovery is thus a fact about the subsystem's inference, not about the world's reversibility. □

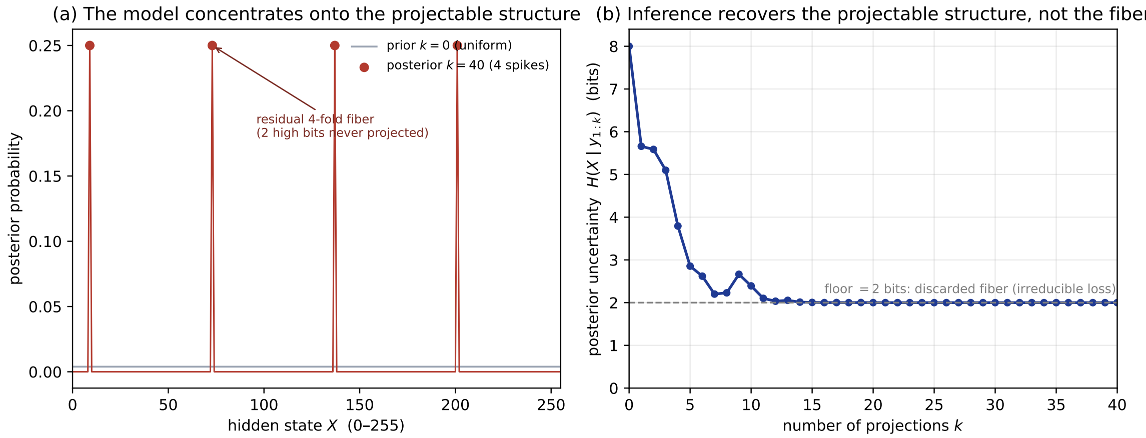

Figure 1: The model recovers the projectable structure by inference, not the discarded fiber. A hidden 8-bit state X has a uniform prior; the projection reveals only its low 6 bits, and those noisily (per-bit flip probability 0.15); the high 2 bits are never projected — the discarded fiber. A Bayesian subsystem updates its posterior on each noisy projection. **(a)** The posterior collapses from uniform (prior, $k = 0$) onto four equal spikes after $k = 40$ projections — the four states sharing the recovered low 6 bits — with an irreducible four-fold ambiguity, the two unprojected high bits. **(b)** The posterior uncertainty $H(X \mid y_{1:k})$ falls from 8 bits toward a floor of exactly 2 bits and no further: inference recovers the projectable structure but never the discarded fiber. The floor is the residue that seeds the incompleteness of the later bridge articles.

The proposition is the positive half of the asymmetry, and the shift it marks is decisive. Inversion would take the projection and produce the preimage in the world, which is blocked; inference takes the projection and produces, in the subsystem, a hypothesis about the preimage, which is not blocked because it is a different operation in a different place. The model does not undo the projection; it guesses what the projection was a projection of, and checks the guess against the next projection. This is how a bounded system inside a lossy world nonetheless comes to track that world: not by recovering the lost bits, but by maintaining a surrogate good enough to predict what comes next. The prior work of the programme argues that bounded systems must do exactly this — that representational isolation forces an internal model and that predictive processing is its evolutionarily inevitable form [3–5]; the present proposition places that necessity as the positive half of the central asymmetry.

5.2 Counterarguments and Responses

Objection 1 (Bayesian updating increases information about the source, violating the data-processing inequality). *Response.* It does not. The data-processing inequality bounds the mutual information between the source and the *signal*, and no amount of posterior updating changes that bound; what updating changes is the subsystem's posterior *given its model and the signals received*, which is a different quantity. Across many projections the subsystem accumulates information from many signals, each within its own bound, and combines them by inference; no single channel's bound is exceeded, and no information about the source is conjured from a signal that did not carry it. The posterior concentrates because many bounded signals jointly constrain the preimage, not because one signal was over-read.

Objection 2 (the model is just a disguised inversion, so the asymmetry is verbal). *Response.* It is not an inversion, and the difference is operational, not verbal. An inversion would map a projection to its preimage as a function, which the many-to-one collapse forbids (Proposition 3.1); the model instead maps a *sequence* of projections to a *posterior distribution* over preimages, which is defeasible, never certain, and bounded below by the fiber the projections never distinguish (Figure 1). Inversion would output a preimage; the model outputs a hypothesis with residual uncertainty. The asymmetry is the difference between a function that does not exist and an inference that does.

5.3 Section Result and Implications

The model recovers the projectable structure by inference: a subsystem maintains a defeasible surrogate of the preimage, conditioned on received projections, concentrating on the consistent preimages without inverting the projection or violating any block (Proposition 5.1, Figure 1). The implication is that recovery and loss are operations of different kinds in different places. We make the asymmetry precise.

6 The Asymmetry Without Paradox

6.1 Two operations, two places

The central asymmetry can now be stated without residue of paradox.

Proposition 6.1 (The asymmetry, and its consistency with irreversibility)**.** *Projection and modeling are distinct operations residing in distinct systems, and they are not inverse to one another:*

- Projection *maps the substrate to a projected value, is many-to-one, physical, and irreversible, and resides in the world (Propositions 3.1–4.1).*
- Modeling *maps a sequence of projections to a posterior over preimages, is defeasible, costly, and fallible, and resides in the subsystem (Proposition 5.1).*

Modeling does not invert projection: it recreates no discarded distinction, so it violates no irreversibility, and it leaves an irreducible floor — the part of the fiber that no projection ever distinguishes, on which the posterior cannot concentrate. The floor is *the residual uncertainty* $H_\infty > 0$ *that persists as projections accumulate without bound whenever the projections share a common collapsed distinction.*

Proof. Non-inversion and consistency with irreversibility are Proposition 5.1: the model lives in the subsystem's states, recreates no discarded bit, and so leaves Proposition 4.1 intact. The floor is the intersection of the fibers the projections fail to distinguish: if every projection is constant on a non-trivial set of preimages — as the high bits are never revealed in Figure 1 — then the posterior is constant on that set too, and its entropy cannot fall below the log-size of the set, giving $H_\infty > 0$ [1, 2]. The model recovers everything the projections jointly distinguish and nothing they jointly omit. □

The asymmetry resolves the paradox that motivated the article. If projection loses irreversibly, how is knowledge possible? The answer is that knowledge is not the world's recovery of what it lost — which is blocked — but the subsystem's inference of what the world is like, which is a different thing in a different place. The subsystem never gets the discarded fiber back; it builds a surrogate that predicts the next projection, and the surrogate's success is its knowledge. The irreversibility of the world is untouched, because nothing the world discarded is recreated; the fallibility of the subsystem is permanent, because the floor never closes. Loss and recovery coexist precisely because they are not inverse: one is a physical map that lost a fiber, the other an inference that hypothesizes it.

6.2 The floor as the seed of incompleteness

The floor H_∞ deserves emphasis, for it is where the bridge points. It is the residual uncertainty about the preimage that no accumulation of projections removes, the part of the substrate that the subsystem's projections never distinguish and its model therefore never resolves. In Figure 1 it is the two high bits, two clean bits of permanent ignorance. In general it is whatever the available projections hold in common as a collapsed distinction — the blind spot of a given receiver's whole projective access. That such a floor exists for any proper projection is the seed of the incompleteness of measurement that a later bridge article develops: a subsystem confined to projections has structure it can never resolve, not for want of effort but because its access is projective and projective access has fibers. The model recovers up to the floor and stops; the floor is the measure of what modeling, however ideal, leaves undone.

6.3 Counterarguments and Responses

Objection 1 (with richer projections the floor could always be lowered to zero, so it is not fundamental). *Response.* A richer set of projections lowers the floor by distinguishing more of the fiber, and in the limit of projections that jointly distinguish everything the floor reaches zero — but that limit is the bijective, non-lossy reduction, the maximal observer at which projection ceases to be projection (Article 10).

For any *proper* projective access — any receiver that genuinely reduces — some distinction is held in common and the floor is positive. The floor is fundamental to projective access as such, not to any particular poverty of it; only the non-projecting limit removes it, and that limit is not available to a subsystem inside the world.

Objection 2 (the asymmetry just relabels the familiar fact that inference is uncertain). *Response.* It locates that fact and explains its necessity. Inference is uncertain not as a contingent matter of insufficient data but because the world it infers about is accessed through lossy projection, whose fibers no inference resolves; the floor is the exact form of the uncertainty, equal to the jointly collapsed distinction. The asymmetry is not the platitude that inference is fallible but the claim that its fallibility is the shadow of the world's irreversible loss, measured by the floor — which is why the same three axes that forbid inversion also set the limit of inference.

6.4 Section Result and Implications

The asymmetry is precise and consistent: projection and modeling are distinct, non-inverse operations in distinct systems, recovery violates no irreversibility, and an irreducible floor remains (Proposition 6.1). The floor is the seed of the incompleteness the bridge will develop. We set the programme of Part III.

7 What the Bridge Opens

The central asymmetry is the engine of the bridge, and the remaining articles of Part III are its consequences for the model-building subsystem. We sketch the programme the asymmetry sets, without discharging it here.

The *arrow of time* is the asymmetry read on the macrostate axis through the model-builder. Article 8 located the second law as the lossiness of the coarse-graining and deferred the direction of the arrow to Part III; the asymmetry supplies the frame in which to take it up. The model-builder records its past as model-states laid down by past projections and faces its future as projections not yet received; the difference between a past that has been projected into the model and a future that has not is the difference between the recovered and the still-to-be-inferred, and the arrow is the gradient of the model's accumulation. Why the initial macrostate was low-entropy — the boundary condition Article 8 flagged — remains a cosmological question, but the experienced arrow is the model-builder's, and the asymmetry is its setting.

The *relocation of the simulation* is the asymmetry read to its limit. A model elaborate enough to predict its own projections in detail becomes a simulation — a surrogate so rich that running it reproduces the stream of projections the subsystem would receive. The asymmetry locates such a simulation precisely: it lives in the subsystem, as the most elaborate form of the recovering model, not in the world as a rival substrate. This dissolves the simulation hypothesis in the form Article 4 touched philosophically:

the question is not whether the world is a simulation but that simulation is what a sufficiently rich model-building subsystem does, and the simulation is the model's, not the substrate's. The later article develops the relocation.

The *incompleteness of measurement* is the floor made into a theorem. The residual H_∞ of Proposition 6.1 is, for any proper projective access, a positive measure of permanent unresolvability; a later article develops this into the statement that a subsystem confined to projections necessarily has structure it cannot measure, and relates it to the apophatic boundary of Article 4 — the difference between what is ill-typed to ask and what is well-typed but forever unresolved. The floor is well-typed and unresolved; the boundary is ill-typed; the bridge keeps them distinct.

These are the model-builder's situation: an arrow it experiences, a simulation it may become, a floor it cannot cross. Each is the central asymmetry — projection loses, the model recovers — read into one facet of the subsystem's life. The present article has established the asymmetry; the rest of Part III lives inside it.

7.1 Counterarguments and Responses

Objection 1 (sketching the programme without discharging it is empty promising). *Response.* The sketches are not arguments and are not offered as such; they state how each later topic instantiates the asymmetry just proved, so that the opening article earns its place by showing what it opens. The arguments are reserved for the articles named, and the epistemic ledger marks these forward-looking claims as deferred. An opening article that did not indicate the part it opens would be the more deficient.

Objection 2 (calling a rich model a "simulation" courts the reification the volume forbids). *Response.* It is the opposite of reification: the relocation insists that the simulation is the *subsystem's model*, not the substrate, precisely to deny that the world is a simulation. Where the reification of Article 2 mistook a projection for the substrate, the relocation mistakes nothing — it places the simulation where it belongs, inside the model-builder, and leaves the substrate axis-free. The later article makes this exact; the sketch only flags it.

7.2 Section Result and Implications

The central asymmetry sets the programme of Part III: the arrow of time, the relocation of the simulation, and the incompleteness of measurement are the model-builder's situation, each a reading of projection-loses-model-recovers (Section 7). The implication is that the bridge is one idea developed in three directions, and the following articles develop them.

8 Epistemic Status: Proved, Imported, Affirmed, Deferred

In the manner of the series, we close with an explicit ledger. *Proved (here)* marks results established in this article; *Imported* marks results from prior articles; *Affirmed (standard)* marks established work the article locates; *Deferred* marks claims set for a later article of Part III.

Claim	Status	Locus
DPI, uncomputability of inversion, second law (axis results)	Imported	Kriger, IST 5 [7], Kriger, IST 6 [8], Kriger, IST 8 [10]
Projection is many-to-one; the three losses are one fiber	Proved (here)	Prop. 3.1
Physical inversion is blocked (DPI, Landauer, second law)	Proved (here)	Prop. 4.1
Model recovery is inference, not inversion	Proved (here)	Prop. 5.1
Recovery violates no irreversibility; the residual floor $H_\infty > 0$	Proved (here)	Prop. 6.1
Bounded systems must model (representational isolation; predictive processing)	Imported (prior corpus)	Kriger [3, 4]
Data-processing inequality; Bayesian consistency	Affirmed (standard)	Cover and Thomas [1], Jaynes [2]
Landauer's principle (un-erasure has no free source)	Affirmed (standard)	Landauer [12]
Arrow of time as the model-builder's gradient	Deferred	Part III
Relocation of the simulation into the model	Deferred	Part III
Incompleteness of measurement (floor as theorem)	Deferred	Part III

The ledger marks the division of labor. The article *imports* the three axis-losses and the prior corpus on why bounded systems must model; it *proves* the unification of the losses, the block on inversion, the inferential character of recovery, and the residual floor; it *affirms* the standard results it applies; and it *defers* the three developments of the asymmetry to the articles of Part III they belong to.

9 Discussion

The result of this article is the central asymmetry of the volume, and the opening of its bridge. The three losses Part II left standing — the data-processing inequality, the uncomputability of inversion, the second law — are one fact: projection is many-to-one, and the fiber it collapses does not uncollapse. Physical inversion is blocked, by the same three axes in concert, so the world does not un-project. And yet structure projection discarded comes to be known, because a model-building subsystem recovers it — not by inverting the projection, which is blocked, but by inference, maintaining a defeasible surrogate of the preimage conditioned on the projections it receives and tested against those still to come. Loss is physical and resides in the world; recovery is inferential and resides in the subsystem; the two are not inverse, and their distinctness is why both are real at once.

The asymmetry earns its place by what it resolves and what it opens. It resolves the paradox of knowledge under lossy projection: knowledge is the model's inference, not the world's inversion, and inference constrains the preimage across many projections without undoing one. It respects irreversibility exactly, because the model recreates no discarded bit but only hypothesizes it, at a cost and with a permanent fallibility. And it seeds the incompleteness the bridge develops, in the irreducible floor that no inference crosses — the part of the fiber every projection holds in common. The prior corpus of the programme had argued, on evolutionary and information-theoretic grounds, that bounded systems must build internal models and that they perceive through representational isolation rather than direct access [3–5]; the present article places that necessity as the positive half of the asymmetry, the half that answers the negative half's loss. The model-builder is not an optional elaboration of the substrate picture but its required complement: where projection loses, only a model recovers, and only by inference.

9.1 Limitations

Three limitations are material. *First*, the inferential recovery is demonstrated for a Bayesian subsystem with a correct generative model; a subsystem with a misspecified model recovers a biased surrogate, and the article does not treat the rich question of model error beyond noting the surrogate's fallibility. *Second*, the floor H_∞ is derived for projections that share a common collapsed distinction; characterizing the floor for arbitrary families of projections, and relating it precisely to the apophatic boundary of Article 4, is the work of the later incompleteness article, not of this one. *Third*, the three developments of the asymmetry — arrow, simulation, incompleteness — are sketched as the programme of Part III and not argued here; their status is deferred, and the sketches carry no probative weight.

9.2 Future Directions

The immediate sequels are the three articles the asymmetry opens: the arrow of time as the model-builder's gradient of accumulation, the relocation of the simulation into the model, and the incompleteness of measurement as the floor made into a theorem. Beyond these, the asymmetry invites a precise question continuous with the completeness conjecture of Article 3: whether the floor H_∞ of a family of projections is itself a projection-theoretic invariant — a property of the family's joint fiber computable from the inter-axis morphisms — so that the limit of what any model can recover would be read off the same diagram that classifies the axes. That question would tie the bridge back to the classification, and is left to the work that can address it.

Acknowledgments

The author thanks colleagues of the Information Physics Institute and the Institute of Integrative and Interdisciplinary Research for sustained discussion of the foundations of the series.

References

[1] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Wiley-Interscience, Hoboken, NJ, 2nd edition, 2006.

[2] E. T. Jaynes. *Probability Theory: The Logic of Science*. Cambridge University Press, Cambridge, 2003.

[3] B. Kriger. Evolutionary and information-theoretic argument for the necessity of representational isolation: why direct perception was never an option for complex systems. IIIR. Zenodo, 2026. doi:10.5281/zenodo.18331202.

[4] B. Kriger. The evolutionary inevitability of predictive processing: a physical constraint argument. IIIR. Zenodo, 2026. doi:10.5281/zenodo.18444910.

[5] B. Kriger. The structural distortion principle: a closed-loop model of perception, attention, and world-maintenance in bounded cognitive systems. IIIR. Zenodo, 2026. doi:10.5281/zenodo.18452700.

[6] B. Kriger. The projection theorem: how known information measures descend from the substrate. IST, Volume VII, Article 3. Zenodo, 2026.

[7] B. Kriger. The axis of the channel: Shannon information as a substrate projection. IST, Volume VII, Article 5. Zenodo, 2026.

[8] B. Kriger. The axis of description: Kolmogorov complexity and the algorithmic projection. IST, Volume VII, Article 6. Zenodo, 2026.

[9] B. Kriger. The axis of regularity: effective complexity and the structured projection. IST, Volume VII, Article 7. Zenodo, 2026.

[10] B. Kriger. The axis of the macrostate: thermodynamic entropy as a substrate projection. IST, Volume VII, Article 8. Zenodo, 2026.

[11] B. Kriger. The axis of the receiver: Bateson's difference and the pragmatic projection. IST, Volume VII, Article 10. Zenodo, 2026.

[12] R. Landauer. Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5(3):183–191, 1961.

A Code for Figure 1

The following Python script generates Figure 1: a Bayesian subsystem infers a hidden 8-bit state from noisy projections that reveal only its low 6 bits, recovering the projectable structure while the high 2 bits remain an irreducible floor.

```
import numpy as np
rng = np.random.default_rng(11)
N, eps, K, X_true = 256, 0.15, 40, 201        # high 2 bits never projected
xs = np.arange(N); low6 = xs & 0b111111; true_low = X_true & 0b111111

def entropy_bits(logp):
    p = np.exp(logp - logp.max()); p /= p.sum(); p = p[p > 0]
    return -np.sum(p * np.log2(p))

logpost = np.zeros(N); ent = [entropy_bits(logpost)]      # uniform prior: 8 bits
cand = np.stack([(low6 >> b) & 1 for b in range(6)], axis=1)
for k in range(1, K + 1):
    obs = [((true_low >> b) & 1) ^ (1 if rng.random() < eps else 0) for b in range(6)]
    ll = np.where(cand == np.array(obs)[None, :], np.log(1-eps), np.log(eps)).sum(1)
    logpost = logpost + ll; logpost -= logpost.max()
    ent.append(entropy_bits(logpost))         # falls 8 -> floor 2 bits
```

The Arrow of Time:
How the Model-Builder Experiences the Block

The arrow as a subsystem's record gradient, the direction from the low-entropy past, and why the block does not flow

Boris Kriger[1,2]

[1] Information Physics Institute
boris.kriger@informationphysicsinstitute.net

[2] Institute of Integrative and Interdisciplinary Research, Toronto
boriskriger@interdisciplinary-institute.org

ORCID: https://orcid.org/0009-0001-0034-2903

The Information Substrate Theory (IST) — Volume VII, Article 12
Part III: The Bridge

Abstract

The macrostate article located the second law as the lossiness of the coarse-graining and deferred the *direction* of the arrow of time to the bridge; the opening bridge article established the central asymmetry — projection loses, the model recovers by inference — and named the arrow the model-builder's gradient of accumulation. This article develops that name into an account. It begins from a constraint the foundations impose: the substrate is a-temporal, time being one of its axes and not a medium it is in, so temporal predicates do not apply to the substrate and the arrow cannot be a flow of it. The arrow is instead a feature of a model-building subsystem's relation to the block. We locate it in three steps. First, the record asymmetry: a model-builder holds records of past projections and not of future ones, a record being a model-state correlated with what generated it, and the laying of a record is thermodynamically irreversible, so the direction of recording aligns with the entropy gradient. Second, the source of the direction: under time-symmetric micro-dynamics a low-entropy state is an entropy minimum, and coarse-grained entropy rises toward both past and future away from it, as a reversible mixing map demonstrates; the direction of the arrow is therefore fixed not by the dynamics but by the low-entropy boundary condition — the past hypothesis — which the substrate frame takes as a cosmological given and does not derive here. Third, the unification: the thermodynamic, record, causal, and psychological arrows are one fact, the model-builder's monotone accumulation of records along the entropy gradient in a world whose entropy was low in its past. Finally we reconcile the block with the experienced flow: the block does not flow, the "now" is the model's record frontier, and the passage of time is the model updating, an emergent ledger laid down in an a-temporal configuration space.

Keywords: arrow of time; block universe; coarse-graining; past hypothesis; records and memory; Landauer's principle; emergent time; model-building subsystem; foundations of physics.

Contents

1 Introduction

Two earlier articles left the arrow of time as unfinished business. The macrostate article located the second law as the lossiness of the coarse-graining — the coarse-grained entropy rises because reversible micro-dynamics filaments the phase-space density below the coarse-graining's resolution — and was careful to add that this located the *lossiness* but not the *direction*: the micro-dynamics is time-symmetric, so the coarse-grained entropy rises in whichever temporal direction one runs it, and the direction of the arrow was deferred to the bridge [6]. The opening bridge article then established the central asymmetry of the volume — projection loses irreversibly, and a model-building subsystem recovers the projectable structure by inference rather than by inversion — and named the arrow, in passing, the model-builder's gradient of accumulation [8]. This article turns the name into an account.

The account is bounded at the outset by a commitment of the foundations. The substrate is a-temporal: time is one of its axes, a facet along which differentiatedness is read, not a medium the substrate sits in or a flow it undergoes [5]. Temporal predicates — before, after, now, passing — are projection-internal; they apply within the time axis and not to the substrate that the axis projects. It follows immediately that the arrow of time cannot be a property of the substrate or a flow of the block. Whatever the arrow is, it is a feature of something *in* the block that has a before and an after for itself — a subsystem with records, a memory, a present. The arrow is the model-builder's, not the world's.

We locate it in three steps that mirror the structure of the opening bridge article. The *record asymmetry* is the first: a model-builder holds records of the projections it has received and none of those it has not, a record being a model-state correlated with what wrote it, and the writing of a record is thermodynamically irreversible, dissipating at least the Landauer cost, so the direction in which records accumulate is the direction of entropy increase [7, 9]. The recovered-past and inferred-future of the asymmetry [8] are, concretely, the recorded and the unrecorded. The *source of the direction* is the second: under time-symmetric dynamics a low-entropy state is an entropy minimum, and coarse-grained entropy rises toward both its past and its future, as a reversible mixing map shows directly; the arrow points the way it does only because the low-entropy state lies in the model-builder's past — the past hypothesis — a boundary condition the substrate frame takes from cosmology and does not derive here [1, 2]. The *unification* is the third: the thermodynamic, record, causal, and psychological arrows are not four mysteries but one, the model-builder's monotone accumulation of records along the entropy gradient in a world with a low-entropy past.

Finally we reconcile the block with the flow we seem to feel. The block does not flow; there is no moving present in the substrate. But a model-builder embedded in the block lays down records monotonically along the entropy gradient, and the frontier of that record-laying is what it calls "now." The passage of time is the model updating — an emergent ledger written in an a-temporal configuration space, in the sense the programme's ledger-time model developed [3, 4]. The block is still; the ledger grows;

and the growing ledger, not a flowing block, is the arrow.

2 Logical Structure of the Argument

The article proceeds from the a-temporality of the substrate to the unification of the arrows.

Step 1 (Section 3): The substrate is a-temporal.
Time is one axis; temporal predicates do not apply to the substrate; the arrow is not a flow of the block (Proposition 3.1).

Step 2 (Section 4): The record asymmetry.
A model-builder records past projections, not future ones; recording is thermodynamically irreversible and aligns with the entropy gradient (Proposition 4.1).

Step 3 (Section 5): The direction from the low-entropy past.
Under time-symmetric dynamics a low-entropy state is an entropy minimum; the direction is fixed by the boundary condition, not the dynamics (Proposition 5.1, Figure 1).

Step 4 (Section 6): The arrows are one, and the block does not flow.
The thermodynamic, record, causal, and psychological arrows are the model-builder's single record gradient; the "now" is the record frontier and the block does not flow (Proposition 6.1).

Step 5 (Section 7): What remains deferred.
The origin of the low-entropy past is a cosmological question the article flags and does not settle.

Step 6 (Section 8): Epistemic status.
A ledger separates what is proved, imported, affirmed, and deferred.

The dependencies are linear: Step 1 rules out the arrow being the substrate's; Step 2 locates it in the subsystem's records; Step 3 supplies the direction from the boundary condition; Step 4 unifies the arrows and reconciles the block with the flow; Step 5 marks the one thing left to cosmology. The thesis is that the arrow is the model-builder's monotone record gradient, directed by the low-entropy past, and that the block does not flow.

3 The Substrate Is A-temporal

3.1 Time is one axis, not a medium

The first step is a constraint the foundations already fixed, here made explicit for the arrow.

Proposition 3.1 (The arrow is not a property of the substrate)**.** *Time is one of the axes along which the substrate's differentiatedness is read, not a medium the substrate is in (Article 1). Temporal predicates — earlier, later, present, elapsing — are projection-internal, applying within the time axis. The substrate, being prior to any axis, satisfies no temporal predicate: it neither flows nor endures nor passes. Hence the arrow of time is not a property of the substrate and not a flow of the block; it can only be a feature of a subsystem that has a before and an after for itself.*

Proof. That time is one axis and that temporal predicates are projection-internal is established in Article 1, where the substrate is shown to satisfy no axis-predicate, temporal ones included: to apply "passes" or "now" to the substrate is to read a projection's predicate back onto the projected, the category error Article 2 diagnosed and Article 4 marked ill-typed [5]. A flow of time would be a temporal predicate of the world as a whole; the world as substrate has none. Therefore any genuine before/after asymmetry — any arrow — must be predicated of something within the time axis, a subsystem with temporally ordered states, not of the substrate. □

The proposition removes a tempting picture before it can mislead. One imagines the arrow as a property of the universe: time *flows*, and the flow has a direction, and that direction is the arrow. On the substrate reading this picture mislocates the arrow in the substrate, where no temporal predicate holds. The block universe of relativity says the same in its own idiom: the four-dimensional whole does not flow, events simply stand in their temporal relations, and the sense of passage is not a feature of the manifold [10]. The substrate frame generalizes the point: not only does the block not flow, the substrate beneath every axis has no temporal character at all, and the arrow must be sought where before and after are actually instantiated — in a subsystem.

3.2 Counterarguments and Responses

Objection 1 (the second law is a fact about the world, so the arrow is the world's, not a subsystem's). *Response.* The second law is a fact about *coarse-grained* entropy (Article 8), and coarse-graining is a reduction along the macrostate axis, not a feature of the substrate; the entropy that increases is a projection's, and its increase is the lossiness of a coarse-graining a subsystem applies. The law is objective given the coarse-graining, as the macrostate article established, but its temporal asymmetry is read within the time axis, not predicated of the substrate. The world's arrow and the subsystem's are not rivals: the world supplies the entropy gradient, the subsystem reads it as before-and-after.

Objection 2 (saying the block does not flow denies the obvious passage of time). *Response.* It denies that passage is a flow of the block, not that there is anything to the experience of passage. Section 6 locates the experience precisely: it is the model-builder's record frontier advancing along the entropy gradient, a real feature of the subsystem, not an illusion to be explained away. The block does not flow and the model-builder does record; both are true, and the experience of passage is the second, not the first.

3.3 Section Result and Implications

The arrow is not a property of the substrate or a flow of the block, since the substrate satisfies no temporal predicate and time is one of its axes (Proposition 3.1). The implication is that the arrow must be located in a subsystem with a before and an after. We locate it in the subsystem's records.

4 The Record Asymmetry

4.1 Memory is of the past

A model-builder has a before and an after in a definite sense: it holds records of some projections and not others, and the records it holds are of the past.

Proposition 4.1 (The record asymmetry)**.** *A model-builder maintains records — model-states correlated with the projections that wrote them. It holds records of past projections and none of future ones, for three linked reasons:*

(i) (Correlation.) *A record is a state correlated with what generated it; only past projections have generated the subsystem's current states, so only they are recorded.*

(ii) (Irreversibility.) *Writing a record is logically irreversible and so, by Landauer's principle, dissipates at least $k_{\mathrm{B}}T\ln 2$ per bit; the direction in which records are written is the direction of entropy increase [7, 9].*

(iii) (The asymmetry of the bridge.) *The recovered-past and inferred-future of the central asymmetry [8] are the recorded and the unrecorded: the past is what inference has recovered into records, the future what no projection has yet delivered.*

Proof. (i) A record's content is fixed by a correlation between the subsystem's state and an external condition; correlations are established when projections are received and processed, so a subsystem's present records correlate with projections already received — its past — and with nothing not yet received. *(ii)* The formation of a record overwrites a prior state, a logically irreversible operation, which by Landauer carries a minimum dissipation [9]; the cost is paid into the environment along the entropy gradient, so recording proceeds in the direction the macrostate axis calls forward (Article 8). *(iii)* By the central asymmetry, recovery is inference over received projections [8]; the recovered structure is held in records, and the not-yet-received future is, by definition, unrecorded. The three give the same asymmetry: records are of the past. □

The proposition makes the abstract asymmetry of the bridge concrete and physical. Article 11 spoke of a recovered past and an inferred future as a structural feature of inference; here that feature is the ordinary fact that a system remembers backward and not forward, and the fact is grounded in the physics of recording. A record is a correlation, correlations are written by the projections that arrive, and writing dissipates; so a model-builder's memory points the way entropy increases, and its present is the frontier where new records are being laid down. The asymmetry of memory is not a separate puzzle from the second law; it is the second law in the register of records.

4.2 Counterarguments and Responses

Objection 1 (one can record the future by prediction, so memory is not only of the past). *Response.* Prediction produces a model-state about the future, but it is a record of *present inference*, not of the future event, which has written nothing into the subsystem. The asymmetry is between records correlated with what *generated* them — necessarily past — and hypotheses *about* what has not yet generated anything. A prediction is the model's inference (Article 11), held in the present and tested when the future becomes the recorded past; it is not a memory of the future, and its eventual confirmation is the writing of a new past-directed record.

Objection 2 (Landauer's cost is tiny and often unpaid, so it cannot ground the record arrow). *Response.* The magnitude is not the point; the *sign* is. However small, the dissipation of recording is non-negative and is paid into the environment along the entropy gradient, so the operation has a thermodynamic direction, and that direction is the arrow. Real memories dissipate far above the Landauer floor, which only strengthens the alignment; the floor establishes that even an ideal record cannot be written against the gradient for free, which is what fixes the direction in principle.

4.3 Section Result and Implications

The model-builder's records are of the past, by correlation, by the irreversibility of recording, and as the concrete form of the bridge's recovered-past/inferred-future asymmetry (Proposition 4.1). The implication is that the subsystem has a definite before-and-after aligned with the entropy gradient. We ask what fixes the gradient's direction.

5 The Direction Comes from the Low-Entropy Past

5.1 A low-entropy state is an entropy minimum

The record asymmetry aligns the arrow with the entropy gradient, but it does not yet say why the gradient points the way it does. The answer is not in the dynamics.

Proposition 5.1 (The direction is a boundary condition, not a law of motion)**.** *Under time-symmetric, reversible micro-dynamics, a low-entropy macrostate is a local minimum of coarse-grained entropy: the entropy rises in both temporal directions away from it. Hence the dynamics does not select a direction for the arrow. The direction is fixed by the boundary condition that the low-entropy state lies in the model-builder's past — the past hypothesis — and not by any law of motion.*

Proof. Reversible dynamics is symmetric under time reversal, so if coarse-grained entropy increases from a low-entropy state in one temporal direction it increases in the other as well: the state is an entropy minimum, not a one-sided start [2, 6]. Figure 1(a) demonstrates this with a reversible mixing map run forward and backward from a low-entropy ensemble; both branches rise toward equilibrium. A direction for the arrow therefore cannot come from the equations of motion, which are indifferent to it; it comes from the contingent fact that the entropy was low in the past — the past hypothesis — which is a boundary condition on the actual history, not a dynamical law [1]. □

The proposition isolates exactly what the dynamics supplies and what it does not. The dynamics supplies the *gradient* — that entropy is higher away from the minimum — but not its *sign*, since the gradient is two-sided. The sign comes from the boundary condition: in our universe the entropy was extraordinarily low in the past, and it is this contingent initial condition, not Newton's or Schrödinger's time-symmetric laws, that makes the future the high-entropy direction and the past the low. The substrate frame takes this boundary condition as given. Why the early universe was in so low-entropy a state is a cosmological question — one the broader programme addresses elsewhere through the topology and dynamics of the early cosmos — and the present article neither needs nor attempts its answer; it needs only that the condition holds, and given that it holds, the arrow's direction follows.

5.2 Counterarguments and Responses

Objection 1 (this makes the arrow contingent, but the arrow feels necessary). *Response.* The arrow's *existence* given a low-entropy boundary is robust — any model-builder in such a world has a record gradient — but its *direction* is indeed contingent on which end of the history has low entropy, and that contingency is the correct result, not a defect. The feeling of necessity is the model-builder's: situated entirely on the rising branch, it cannot record the other direction and so cannot imagine the arrow reversed.

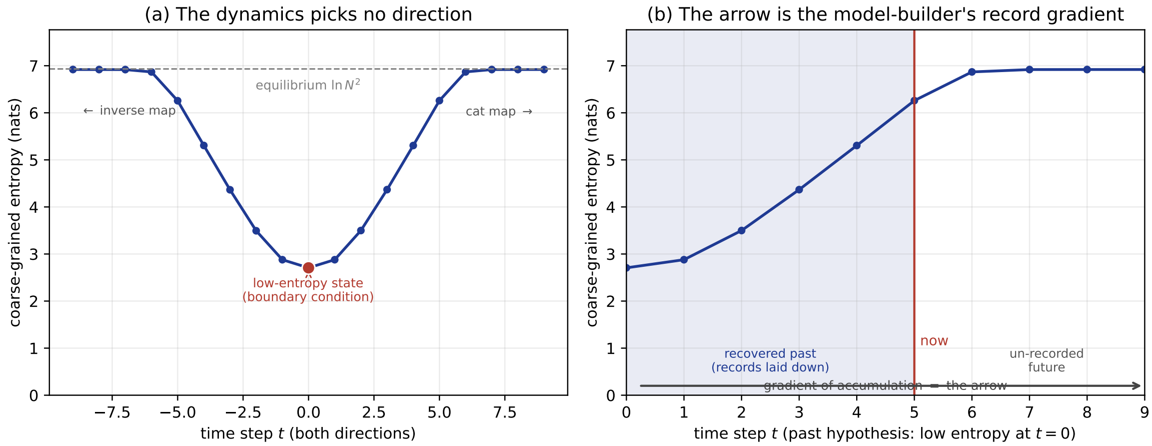

Figure 1: The direction of the arrow is a boundary condition, not a law of motion. **(a)** An ensemble in a low-entropy state at $t = 0$ is evolved forward by Arnold's cat map and backward by its inverse, both reversible and area-preserving. The coarse-grained entropy rises toward *both* past and future: under time-symmetric dynamics a low-entropy state is an entropy minimum, so the dynamics selects no direction. **(b)** Given the past hypothesis — the low-entropy state lies in the model-builder's past — the experienced arrow is the monotone gradient of accumulated records: behind the moving "now" the past is recorded and recovered; ahead the future is unrecorded. The arrow is this gradient of accumulation, and its direction is the direction in which entropy was low.

The necessity is perspectival, the contingency structural, and the figure shows why the two coexist.

Objection 2 (deferring the low-entropy past to cosmology leaves the arrow unexplained). *Response.* It leaves one thing unexplained — the boundary condition — and explains everything else given it. That is the honest division: the article shows that the arrow is the record gradient and that its direction tracks the low-entropy end, which is a complete account modulo the boundary condition; the origin of the boundary condition is a distinct, cosmological question that no account of the arrow's *structure* can settle, and conflating the two has long made the arrow seem more mysterious than it is. The structure is settled here; the boundary condition is flagged and deferred.

5.3 Section Result and Implications

Under time-symmetric dynamics a low-entropy state is an entropy minimum, so the direction of the arrow is fixed by the low-entropy past — a boundary condition — and not by the laws of motion (Proposition 5.1, Figure 1). The implication is that, given the past hypothesis, the model-builder's record gradient has a definite direction. We unify the arrows and reconcile the block with the flow.

6 The Arrows Are One, and the Block Does Not Flow

6.1 One gradient, several names

The arrows of time are usually enumerated — thermodynamic, psychological, causal, radiative, cosmological — and their alignment treated as a coincidence to be explained. On the present account they are not several but one.

Proposition 6.1 (The arrows are one; the block does not flow)**.** *In a world with a low-entropy past, the thermodynamic arrow (entropy increase), the record or psychological arrow (memory of the past), and the causal arrow (control of the future through the present) are the same fact: the model-builder's monotone accumulation of records along the entropy gradient. They align because they are one gradient under several descriptions, not several gradients that happen to agree. Moreover the block does not flow: there is no moving present in the substrate. The "now" is the frontier of the model-builder's record-laying, and the passage of time is the model updating — an emergent ledger written in an a-temporal configuration space.*

Argument. Each named arrow is the record gradient in a register. The thermodynamic arrow is the gradient itself (Article 8). The psychological arrow is the gradient read through records, which are of the past (Proposition 4.1). The causal arrow is the gradient read through intervention: a subsystem can correlate its present records with past states (control of the future requires the present, which is recorded) but cannot write records of the future, so influence runs from the recorded toward the unrecorded. These are the same monotone accumulation, so they align necessarily, not coincidentally. That the block does not flow is Proposition 3.1; that the "now" is the record frontier and passage is the model updating is the ledger-time reading, in which time emerges from the irreversible laying-down of records in an otherwise a-temporal configuration space [3, 4]. The substrate is still; the ledger grows; the growth is the passage. □

The proposition resolves two long-standing puzzles at once. The first is the alignment of the arrows: it is no coincidence that we remember the low-entropy direction, that causes precede effects in that same direction, and that entropy increases in it, because all three are the single gradient of record accumulation, and a subsystem cannot have its memory point one way and its entropy the other. The second is the tension between the block universe and the felt passage of time. The block does not flow, and yet time passes for us; the resolution is that passage is not a property of the block but of the model-builder within it, whose records accumulate monotonically and whose frontier of accumulation is the present. Nothing in the block moves; the model's ledger lengthens, and the lengthening is what passage is. This is the programme's ledger-time picture, in which an emergent, irreversible record in an a-temporal space is the whole of what time's passage amounts to [3].

6.2 Counterarguments and Responses

Objection 1 (the radiative and cosmological arrows are not obviously the record gradient, so the unification is incomplete). *Response.* The article claims the unification for the thermodynamic, psychological, and causal arrows, which are the model-builder's directly, and treats the radiative arrow (retarded over advanced radiation) as a further consequence of the same low-entropy past, since coherent outgoing radiation into a low-entropy environment is the field-theoretic face of the entropy gradient. The cosmological arrow — the direction of cosmic expansion — is not claimed to be the record gradient but is the source of the low-entropy past that orients it; it is the boundary condition's origin, deferred to cosmology (Section 7), not another arrow to be unified.

Objection 2 (ledger-time reintroduces a flowing present the article just denied). *Response.* It does not: the ledger grows in the sense that more records stand in the block at later times, not in the sense that a present sweeps along a track. There is no extra fact of "which now is lit up"; there is only, at each time, a model-builder whose records extend up to that time and not beyond, and "now" is indexical, naming each such frontier from within. The block contains all the frontiers at once, none privileged; passage is the relation each frontier bears to its own records, not a spotlight moving over the block. The ledger metaphor is of accumulation across the block, not of a flow through it.

6.3 Section Result and Implications

The thermodynamic, psychological, and causal arrows are one gradient of record accumulation, aligned necessarily; the block does not flow, the "now" is the record frontier, and passage is the model updating (Proposition 6.1). The implication is that the arrow is fully located as a feature of the model-builder, modulo the one boundary condition. We mark that condition's status.

7 What Remains Deferred

One thing the article does not do is explain the low-entropy past. The direction of the arrow tracks the low-entropy end of the history (Proposition 5.1); that the early universe was in an extraordinarily low-entropy state is the boundary condition on which the whole direction turns; and why it was so is a cosmological question, not one any account of the arrow's structure can answer. The substrate frame is explicit about the division. The *structure* of the arrow — that it is the model-builder's record gradient, that its several names are one, that the block does not flow — is settled here, and settled without appeal to the boundary condition's origin. The *direction*, given that the entropy was low at one end, is settled too. Only the *origin* of the low-entropy end is deferred, and it is deferred to cosmology, where the broader programme addresses the early universe's state through the topology and dynamics of the cosmos [5]. Conflating these three — treating the unexplained boundary condition as if it left the arrow's structure or direction unexplained — is what has made the arrow seem more obscure than it is. The article keeps them apart: structure and direction accounted for, origin flagged.

This deferral is of a piece with the volume's discipline. The macrostate article deferred the direction to the bridge; the bridge has now supplied it, modulo the boundary condition; and the boundary condition's origin is deferred, in turn, to the cosmology the programme's other volumes pursue. Each deferral is honest and bounded, and none is a gap in the present account, which concerns the arrow as the model-builder experiences it.

8 Epistemic Status: Proved, Imported, Affirmed, Deferred

In the manner of the series, we close with an explicit ledger. *Proved (here)* marks results established in this article; *Imported* marks results from prior articles; *Affirmed (standard)* marks established work the article locates; *Deferred* marks the boundary-condition question handed to cosmology.

Claim	**Status**	**Locus**
Substrate a-temporal; time is one axis	Imported	Kriger, IST 1 [5]
Second law as lossiness of coarse-graining (direction deferred)	Imported	Kriger, IST 8 [6]
Central asymmetry: projection loses, model recovers	Imported	Kriger, IST 11 [8]
Arrow is not a property of the substrate / flow of the block	Proved (here)	Prop. 3.1
Record asymmetry: memory is of the past	Proved (here)	Prop. 4.1
Low-entropy state is an entropy minimum (direction is a boundary condition)	Proved (here)	Prop. 5.1
Thermodynamic, record, causal arrows are one gradient	Proved (here)	Prop. 6.1
Block does not flow; "now" is the record frontier (ledger-time)	Proved (here)	Prop. 6.1; Kriger [3]
Recording is irreversible (Landauer cost, sign not magnitude)	Affirmed (standard)	Landauer [9]
Time-symmetry of micro-dynamics; the past hypothesis	Affirmed (standard)	Albert [1], Boltzmann [2], Price [10]
Origin of the low-entropy past	Deferred	cosmology

The ledger marks the division of labor. The article *imports* the a-temporal substrate, the second law, and the central asymmetry; it *proves* that the arrow is the model-builder's record gradient, that the direction is a boundary condition, that the arrows are one, and that the block does not flow; it *affirms* the standard results it applies; and it *defers* the origin of the low-entropy past to cosmology.

9 Discussion

The result of this article is an account of the arrow of time within the substrate frame. The arrow is not a property of the substrate, which is a-temporal, nor a flow of the block, which is still; it is a feature of a model-building subsystem embedded in the block. Concretely it is the subsystem's monotone gradient of accumulated records: a record is of the past because it correlates with what wrote it and because writing dissipates along the entropy gradient, so memory points the way entropy increases. The direction that gradient takes is not selected by the time-symmetric micro-dynamics, for which a low-entropy state is a two-sided entropy minimum, but by the low-entropy boundary condition — the past hypothesis — which the substrate frame takes from cosmology. Given that condition, the thermodynamic, psychological, and causal arrows are one gradient under several names, aligned necessarily; and the felt passage of time, far from being a flow of the block, is the model-builder's ledger lengthening across a block that does not move.

The account fits the architecture the bridge is building. Article 11 established that projection loses and the model recovers by inference, with a recovered past and an inferred future; this article shows that the recovered past is the recorded past and the inferred future the unrecorded, so the bridge's central asymmetry, read in the register of records, *is* the arrow of time. The programme's earlier work prepared the reading: the ledger-time model derived time's passage from the irreducible laying-down of records in an a-temporal configuration space, and the work on atemporal memory argued that storage is timeless and that time belongs to retrieval, not to the stored [3, 4]. The present article places those results as the bridge's account of the arrow: the substrate is timeless, the records accumulate, and the accumulation, directed by the low-entropy past, is everything the arrow of time amounts to for the subsystem that has one.

9.1 Limitations

Three limitations are material. *First*, the article accounts for the arrow's structure and direction but defers the origin of the low-entropy past to cosmology; a reader who wants the arrow explained *all the way down*, boundary condition included, will find that question marked and not answered here. *Second*, the unification is claimed for the thermodynamic, psychological, and causal arrows; the radiative arrow is treated as a consequence of the same boundary condition and the cosmological arrow as the boundary condition's source, but neither is developed in detail. *Third*, the ledger-time reconciliation of the block with passage is stated and grounded in the prior corpus but not formalized here; its precise formulation, in which records are the irreversible events of an a-temporal configuration space, belongs to the work cited and to the incompleteness article that follows.

9.2 Future Directions

The immediate sequel is the incompleteness article, which takes up the residual floor of Article 11 — the part of the fiber no projection resolves — and makes it a theorem on the limits of measurement, relating it to the apophatic boundary of Article 4. The arrow connects to that work through the record frontier: what a subsystem can ever record is bounded by what its projections distinguish, so the floor bounds memory as it bounds inference. Beyond the volume, the account invites a precise question continuous with the completeness conjecture of Article 3: whether the record gradient can be made a projection-theoretic quantity — a monotone on the macrostate axis whose increase along the morphism to the channel axis is the rate of record accumulation — so that the arrow's strength would be read off the same diagram that classifies the axes. That question is left to the work that can address it; the present article's task was to locate the arrow as the model-builder's, and the arrow is located.

Acknowledgments

The author thanks colleagues of the Information Physics Institute and the Institute of Integrative and Interdisciplinary Research for sustained discussion of the foundations of the series.

References

[1] D. Z. Albert. *Time and Chance*. Harvard University Press, Cambridge, MA, 2000.

[2] L. Boltzmann. Über die Beziehung zwischen dem zweiten Hauptsatze der mechanischen Wärmetheorie und der Wahrscheinlichkeitsrechnung respektive den Sätzen über das Wärmegleichgewicht. *Sitzungsberichte der Kaiserlichen Akademie der Wissenschaften, Wien*, 76:373–435, 1877.

[3] B. Kriger. The ledger time model: emergent time from irreversible entanglement validation in an atemporal quantum configuration space. IIIR. Zenodo, 2026. doi:10.5281/zenodo.18263726.

[4] B. Kriger. Atemporality of mental memory space: a structural hypothesis grounded in resource constraints, cyclical closure, and reconstructive retrieval. IIIR. Zenodo, 2026. doi:10.5281/zenodo.18381912.

[5] B. Kriger. The information substrate: differentiation as the ground of information. IST, Volume VII, Article 1. Zenodo, 2026.

[6] B. Kriger. The axis of the macrostate: thermodynamic entropy as a substrate projection. IST, Volume VII, Article 8. Zenodo, 2026.

[7] B. Kriger. The axis of cost: Landauer's principle and the infodynamics of DESI DR1. IST, Volume VII, Article 9. Zenodo, 2026.

[8] B. Kriger. Projection loses, the model recovers: the central asymmetry and the opening of the bridge. IST, Volume VII, Article 11. Zenodo, 2026.

[9] R. Landauer. Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5(3):183–191, 1961.

[10] H. Price. *Time's Arrow and Archimedes' Point: New Directions for the Physics of Time*. Oxford University Press, New York, 1996.

A Code for Figure 1

The following Python script generates Figure 1: a low-entropy ensemble evolved forward by Arnold's cat map and backward by its inverse, with coarse-grained entropy rising in both temporal directions.

```
import numpy as np
rng = np.random.default_rng(12); M, N, T = 40000, 32, 9
x0 = 0.45 + 0.10*rng.random(M); y0 = 0.45 + 0.10*rng.random(M)

def cat(x, y):   return (2*x + y) % 1.0, (x + y) % 1.0        # forward
def icat(x, y):  return (x - y) % 1.0, (-x + 2*y) % 1.0       # inverse (reverse)

def coarse_entropy(x, y, N):
    H, _, _ = np.histogram2d(x, y, bins=N, range=[[0,1],[0,1]])
    p = H.ravel(); p = p[p > 0] / p.sum()
    return -np.sum(p * np.log(p))

# forward branch (cat) and backward branch (icat) from the SAME low-entropy state:
# coarse-grained entropy rises in BOTH directions -> minimum at t=0.
```

The Relocation of the Simulation: Why Simulation Lives in the Model, Not in the World

The simulation as the limit of the recovering model, the ill-typing of the world-as-simulation question, and the halting of the regress

Boris Kriger[1,2]

[1] Information Physics Institute
boris.kriger@informationphysicsinstitute.net

[2] Institute of Integrative and Interdisciplinary Research, Toronto
boriskriger@interdisciplinary-institute.org

ORCID: https://orcid.org/0009-0001-0034-2903

The Information Substrate Theory (IST) — Volume VII, Article 13
Part III: The Bridge

Abstract

The opening bridge article noted that a model elaborate enough to predict its own projection-stream becomes a simulation, and that such a simulation lives in the subsystem, not in the world. This article develops the remark into an account and applies it to the simulation hypothesis. The hypothesis is usually posed as a question about the world: is the world the output of a computation running on a deeper substrate? We argue that this framing mislocates a projection. The word "simulation" has two readings, and only one is well-typed. On the first reading, a simulation is a computational substrate beneath physics of which our world is the output; this reifies a projection — the algorithmic description of a system — into a second substrate, and positing a substrate beneath the substrate is the category error the foundations rule ill-typed, not false. There is no well-formed answer to "is the world running on a deeper computer," because the question reads a projection's predicate back onto the substrate. On the second reading, a simulation is what a sufficiently rich model-building subsystem does: a model run forward to reproduce the projections the subsystem would receive. This reading is well-typed and real, and it locates the simulation in the subsystem as the limit of the recovering model — a map that has become executable but is still a map, not the territory. We dissolve the regress of the simulation argument: a tower of nested simulations is a tower of nested *models* inside one a-temporal substrate, not a stack of substrates, and the regress halts because each level is a subsystem in the one substrate, not a substrate of its own. Finally we locate the grain of truth in the hypothesis — representational isolation, the fact that a subsystem meets the world only through its model — as a fact about the receiver axis, not evidence of a deeper substrate. The article claims neither that the world is a simulation nor that it is not; it shows that the question, as posed, is mislocated.

Keywords: simulation hypothesis; substrate and projection; ill-typed questions; model-building subsystem; representational isolation; map and territory; nested simulations; foundations of physics.

Contents

1 Introduction

The opening article of the bridge ended with a remark it did not develop. A model rich enough to predict its own projection-stream in detail, it observed, becomes a simulation — a surrogate so complete that running it reproduces the projections the subsystem would receive — and the central asymmetry locates such a simulation in the subsystem, as the most elaborate form of the recovering model, rather than in the world as a rival substrate [10]. This article develops the remark into an account and turns it on the simulation hypothesis, the proposal that the world we inhabit might itself be a simulation [1].

The hypothesis is almost always posed as a question about the world. Is what we take for physical reality the output of a computation running on a machine at some deeper level — a machine in a "base reality" of which our universe is a rendered approximation? Posed so, the question invites a probabilistic argument: if simulations are feasible and civilizations run many of them, then simulated observers vastly outnumber unsimulated ones, and we should expect to be among the simulated. The argument has force only if its central noun is well-formed — only if "the world is a simulation" picks out a genuine state of affairs the world might be in. We argue that it does not, and that seeing why is a clean application of the foundations laid in Part I.

The word "simulation" carries two readings, and the work of the article is to separate them. On the first, a simulation is a *computational substrate beneath physics*: a deeper level of reality, running our world as a program. This reading reifies a projection. The algorithmic, computational description of a system is a projection along the description axis (Article 6); to posit it as a substrate underlying the world is to mistake a projection for the projected, the error Article 2 diagnosed, and to posit a substrate *beneath* the substrate is to posit something beyond the block, which Article 4 marked ill-typed — non-denoting, neither true nor false [5, 6]. "Is the world running on a deeper computer" has, on this reading, no well-formed answer, because it reads a projection's predicate back onto the axis-free substrate.

On the second reading, a simulation is *what a sufficiently rich model-building subsystem does*: a model, the kind the receiver axis terminates in (Article 10) and the bridge showed to recover structure by inference (Article 11), elaborated until running it forward reproduces the subsystem's projection-stream. This reading is well-typed and real. It locates the simulation in the subsystem, as the limit of the recovering model, defeasible and costly and never the substrate — an executable map, which for all its detail is still a map and not the territory. The relocation is the article's central move: the simulation, real on the second reading, is moved from the world (where the first reading wrongly put it) to the subsystem (where it belongs).

With the relocation in hand, the regress dissolves and a grain of truth is located. The tower of nested simulations that powers the probabilistic argument is, on the substrate frame, a tower of nested *models* inside one a-temporal substrate, not a stack of substrates; the regress halts because each level is a subsystem in the one substrate, not a substrate of its own, and the count of levels is a count of models, telling us nothing about whether "the world is a simulation," which is the ill-typed question. And the grain of truth the hypothesis feeds on — the sense in which we do seem to live "inside" something — is representational isolation: a subsystem meets the world only through its model [2, 3], so the world it experiences is already its model. That is a fact about the receiver axis, not evidence of a deeper computer. The article claims neither that the world is a simulation nor that it is not; it shows the question, as posed, to be mislocated.

2 Logical Structure of the Argument

The article separates two readings and follows each to its place.

Step 1 (Section 3): Two readings of "simulation."
The world-as-substrate reading and the model-in-subsystem reading are distinguished (Proposition 3.1).

Step 2 (Section 4): The world-as-simulation reading is ill-typed.
It reifies a projection into a second substrate and posits something beyond the block; ill-typed, not false (Proposition 4.1).

Step 3 (Section 5): The simulation is the limit of the recovering model.
On the well-typed reading the simulation is a model run forward, located in the subsystem; an executable map is still a map (Proposition 5.1, Figure 1).

Step 4 (Section 6): The regress halts.
Nested simulations are nested models in one substrate, not a tower of substrates; the probabilistic argument inherits the ill-typing (Proposition 6.1).

Step 5 (Section 7): The grain of truth.
Representational isolation is the receiver-axis fact the hypothesis feeds on, not a deeper substrate (Proposition 7.1).

Step 6 (Section 8): Epistemic status.
A ledger separates what is proved, imported, and affirmed, and states what the article does not claim.

The dependencies are linear: Step 1 disambiguates; Steps 2 and 3 dispatch the two readings, one as ill-typed, the other as a located reality; Step 4 dissolves the regress that the first reading powered; Step 5 explains the hypothesis's appeal without conceding it. The thesis is that "simulation" names a real activity of subsystems and not a possible substrate of the world, and that the simulation hypothesis, posed as a question about the world, is mislocated.

3 Two Readings of "Simulation"

3.1 The substrate reading and the subsystem reading

The disambiguation is the whole foundation of the article, and it is worth stating sharply.

Proposition 3.1 (Two readings of "simulation")**.** *The term "simulation" admits two distinct readings:*

(A) (Substrate reading.) *A simulation is a computational substrate beneath the world: a deeper level of reality on which the world runs as a program, of which the world is the output.*

(B) (Subsystem reading.) *A simulation is the activity of a model-building subsystem: a model elaborated until running it forward reproduces the projections the subsystem would receive.*

The two are not variants of one notion. Reading (A) places the simulation below *the world, as its ground; reading (B) places it* within *a subsystem of the world, as a model. The simulation hypothesis trades on reading (A); the bridge's recovering model is reading (B).*

Proof. That the readings are distinct follows from where each puts the simulation. Reading (A) makes the simulation ontologically prior to the world — the world depends on it, as an output on its program — so the simulation is a substrate. Reading (B) makes the simulation ontologically posterior to and contained in a subsystem — the model depends on the subsystem that maintains it — so the simulation is a projection within the world (Articles 10, 11). Priority and posteriority to the world cannot coincide, so the readings are genuinely two. The simulation hypothesis asks whether the world stands to some computation as output to program, which is reading (A); the bridge's model run forward is reading (B). □

The proposition is easy to state and easy to conflate in practice, because the same vivid image — a world generated by a computation — serves both readings, and the slide from one to the other is silent. When we run a detailed model and watch it reproduce what we would observe, we have a simulation in sense (B), in our subsystem; the hypothesis then asks, of the world itself, whether it is a simulation in sense (A), beneath everything. The two questions feel like one because the image is shared, but they are opposite in their ontology, and the rest of the article turns on keeping them apart.

3.2 Counterarguments and Responses

Objection 1 (the readings are the same: a simulation in a subsystem is itself a substrate for whatever it simulates). *Response.* A simulation in a subsystem is a substrate for the *simulated* entities, in the sense that they depend on it; but it is not a substrate for *the world*, which is what reading (A) requires. The hypothesis asks whether *our* world rests on a deeper computation, not whether the entities inside a model we run rest on the model. The two readings differ exactly on which world is claimed to rest on what: reading (B)'s simulated entities rest on a model in a subsystem; reading (A)'s world rests on a substrate beneath it. The objection collapses the distinction by changing which world is at issue.

Objection 2 (the distinction is merely verbal; physically a simulation is a simulation). *Response.* The distinction is the opposite of verbal: it is ontological, concerning whether the simulation is prior to the world (a substrate) or contained in it (a model). The same physical process — a computation reproducing a stream of observations — is reading (B), a model in a subsystem, whether or not anyone also entertains reading (A); reading (A) adds the metaphysical claim that the world as a whole is such an output, which is a further and, we argue, ill-typed assertion. The physics is all on the side of (B); (A) is an interpretation laid over it.

3.3 Section Result and Implications

"Simulation" has a substrate reading and a subsystem reading, opposite in ontology and silently conflated because they share an image (Proposition 3.1). The implication is that the simulation hypothesis must be evaluated under the substrate reading it actually employs. We evaluate it.

4 The World-as-Simulation Reading Is Ill-Typed

4.1 Reifying a projection into a second substrate

Under reading (A), the simulation hypothesis makes a definite ontological claim, and the foundations have already classified claims of its form.

Proposition 4.1 (The world-as-simulation reading is ill-typed)**.** *The substrate reading posits a computational level beneath the world of which the world is the output. The computational description of a system is a projection along the description axis (Article 6); to posit it as a substrate underlying the world reifies that projection, the error of mistaking the projected for the substrate (Article 2). Moreover it posits a substrate* beneath *the substrate — a level more fundamental than the axis-free ground — which is to posit something beyond the block, ruled ill-typed by the categorial boundary (Article 4).*

Hence "the world is a simulation," on the substrate reading, is ill-typed: non-denoting, neither true nor false, not a state of affairs the world might be in.

Proof. A simulation in sense (A) is a computation, and a computation is specified by an algorithm — a projection along the description axis (Article 6). To say the world *is* such a computation, beneath its physics, is to take that projection as the world's substrate, which is the reification Article 2 diagnosed: the projection is treated as the projected [5]. Further, the posited computational substrate is claimed to be more fundamental than the world, hence prior to the axis-free substrate the foundations identify as the ground; but the substrate admits nothing prior or beneath, being that of which axes are projections, so a "deeper substrate" is a predicate applied beyond the block, which Article 4 marks ill-typed — a question that fails to denote rather than one with a hidden answer [6]. Both diagnoses give the same verdict: the claim is ill-typed. □

The verdict needs care, for ill-typed is not the same as false, and the difference is the article's chief precaution. To say "the world is a simulation" is ill-typed is not to say the world is definitely not a simulation; that would itself presuppose the question well-formed and answer it in the negative. It is to say that the question, on the substrate reading, does not pick out a genuine alternative — that there is no fact of the matter it asks after, because it applies to the substrate a predicate (running-on-a-deeper-computer) that belongs to projections within the world. The honest response to "is the world a simulation?" is therefore not yes and not no but a diagnosis: the question, so posed, mislocates a projection, and once the projection is relocated (Section 5) nothing well-formed remains to answer.

4.2 Counterarguments and Responses

Objection 1 (we can clearly conceive the world being a simulation, so the question is well-formed). *Response.* What is conceivable is reading (B): a world-sized model, run in some subsystem, reproducing a stream of observations — a vivid and coherent picture. Reading (A) borrows that picture's vividness and adds the claim that *our* world, as a whole, is such an output beneath everything; the added claim is what fails to denote, because it predicates of the axis-free substrate a relation (output-of-a-deeper-program) that has sense only within the world. Conceivability of the model does not confer well-formedness on the substrate claim; it is exactly the conflation Proposition 3.1 warned of.

Objection 2 (physics already describes the world computationally, so a computational substrate is not a category error). *Response.* Physics describes the world along the description and other axes, and those descriptions are projections, legitimately so (Part II); the error is not in describing the world computationally but in promoting the computational *description* to the world's *substrate*, beneath its physics. Article 10a was at pains to make the same distinction for the cosmic web: a computational regularity is realized at a scale, but the substrate is not the computation.

Computational description is welcome; a computational substrate beneath the substrate is the ill-typed addition.

4.3 Section Result and Implications

On the substrate reading the simulation hypothesis reifies a projection and posits something beyond the block, and is therefore ill-typed — non-denoting, not false (Proposition 4.1). The implication is that the real content of "simulation" lies in the other reading. We develop it.

5 The Simulation Is the Limit of the Recovering Model

5.1 The executable map

Under reading (B), "simulation" names something real and already in hand from the bridge.

Proposition 5.1 (The simulation is the limit of the recovering model)**.** *A model-building subsystem maintains a surrogate of the preimage, recovering the projectable structure by inference (Article 11). As the surrogate is elaborated to predict the subsystem's projection-stream in increasing detail, it approaches a* simulation*: a model whose forward execution reproduces the projections the subsystem would receive. The simulation is thus the limit of the recovering model. It is located in the subsystem; it is defeasible, costly to run, and tested against further projections; and it is not the substrate. An executable map is still a map: running a model does not convert it into the territory it models.*

Argument. By Article 11 the recovering model is a surrogate maintained in the subsystem, updated to predict projections; a simulation is this surrogate carried to the point where forward execution generates the projection-stream, the most elaborate form of prediction. It inherits the recovering model's location and status: in the subsystem, realized in the subsystem's states, costly by Landauer to run and update, fallible against the next projection (Articles 9, 11). It is not the substrate, because it is a projection within a subsystem (reading B, Proposition 3.1), and a projection is not the projected (Article 2). The map-becomes-executable image makes the point: detail and dynamism do not change a model's type; an executable map predicts the territory's projections without being the territory. □

The proposition completes the relocation. Reading (A) put the simulation beneath the world, where it is ill-typed; reading (B) puts it within a subsystem, where it is a located reality — the limit of the recovering model the bridge has been developing. The vivid

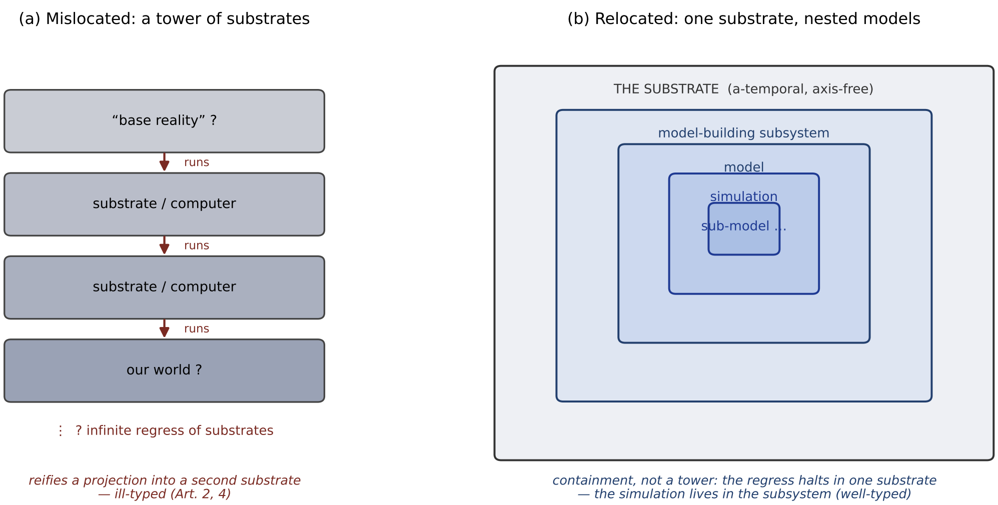

Figure 1: The relocation of the simulation. **(a)** The simulation hypothesis as usually posed: a tower of substrates, each running the level above, regressing without end. On the substrate frame this reifies a projection into a second substrate, and a substrate beneath the substrate is ill-typed (Articles 2, 4). **(b)** The relocation: one a-temporal, axis-free substrate, within which sits a model-building subsystem, within it a model, at the model's limit a simulation, and within the simulation possibly a sub-model. These are nested *models* inside one substrate — containment, not a tower of substrates — so the regress halts, because each level is a subsystem in the one substrate rather than a substrate of its own. The simulation lives in the subsystem; that is well-typed.

scenario that makes the simulation hypothesis gripping — a complete, dynamic, world-sized reproduction of experience — is entirely available on reading (B): a subsystem can in principle build a model so rich that running it yields its whole projection-stream. What reading (B) denies is only that such a model is the world's substrate. It is the subsystem's map, executable and detailed, but a map, and the territory it maps is the projected world, which remains a projection of the one axis-free substrate.

5.2 Counterarguments and Responses

Objection 1 (a perfect simulation is indistinguishable from the world, so the map/territory distinction is empty). *Response.* Indistinguishability of the *projection-stream* does not collapse the type distinction. A simulation perfect in its outputs reproduces every projection a subsystem would receive, and to that subsystem the two are observationally identical — which is precisely the condition of representational isolation (Section 7), not evidence that the map is the territory. The map and territory differ in type, not necessarily in observable output: one is a surrogate maintained in a subsystem, the other the projected world; their agreement in projections is the model's success, not its promotion to substrate. The distinction is between what a thing *is* and what it *outputs*, and only the latter can be made to coincide.

Objection 2 (if the simulation is only ever a map, genuine artificial worlds are impossible). *Response.* Not at all: a subsystem can build models containing entities for which the model is the substrate (reading B's simulated entities), and those entities are as real as any projection within the world. What is denied is only that *our* world is such an entity in a deeper model — reading (A) — which is the ill-typed claim. Artificial worlds, simulated agents, executable models of arbitrary richness are all well-typed realities in subsystems; the relocation affirms them and only refuses to read the actual world as one of them beneath itself.

5.3 Section Result and Implications

On the subsystem reading the simulation is the limit of the recovering model, located in the subsystem, defeasible and costly and never the substrate — an executable map that is still a map (Proposition 5.1, Figure 1). The implication is that the simulation is real but relocated, from the world to the subsystem. We turn the relocation on the regress.

6 The Regress Halts

6.1 Nested models in one substrate

The probabilistic force of the simulation hypothesis comes from a regress: simulations within simulations, a tower of levels, with the inference that we are likely far down it. The relocation dissolves the regress.

Proposition 6.1 (The regress halts: nested models, one substrate)**.** *A simulation in a subsystem may itself contain a model-building sub-subsystem running its own simulation, and so on; this is a regress of nested* models*, each contained in the one before. On the substrate frame it is not a tower of substrates: every level is a subsystem in the one a-temporal substrate, not a substrate of its own (reading B). The regress therefore halts in that single substrate — there is no stack of grounds to descend — and the count of nested levels is a count of models, which says nothing about whether "the world is a simulation," the ill-typed question of reading (A). The probabilistic argument, presupposing a tower of substrates, inherits the ill-typing.*

Proof. By Proposition 5.1 each simulation is a model in a subsystem; a simulation containing a sub-simulation is a model containing a sub-model, by containment, all within the same world and so within the one substrate (Articles 1, 2). There is no point in the nesting at which a new substrate is introduced, since every level is reading (B), a projection within a subsystem, never reading (A), a substrate beneath the world. Hence the regress is bounded below by the single substrate and halts there. The probabilistic argument counts levels of a supposed tower of *substrates* and infers our likely position in it; but the levels are models, not substrates, and "which substrate-level are we?" is the ill-typed question of Proposition 4.1, so the count supports no inference about the world's being a simulation. □

The proposition removes the engine of the simulation argument. That argument pictures reality as a stack: a base reality runs simulations, those run sub-simulations, and since the simulated outnumber the base inhabitants, a random observer is probably simulated. The picture requires that each level be a substrate for the next, a genuine descent of grounds. On the relocation there is no such descent: the nesting is containment of models within models, all inside one substrate, and "how deep are we in the stack of substrates?" has no answer because there is no stack of substrates — only one substrate with models nested inside its subsystems. The vivid arithmetic of simulated versus base observers presupposes exactly the tower the relocation denies, and so it does not get started.

6.2 Counterarguments and Responses

Objection 1 (even nested models give nested observers, so the counting argument survives without a tower of substrates). *Response.* Nested models can contain simulated observers, and one may count them; but the simulation hypothesis's conclusion is about whether *our* world is a simulation in sense (A), beneath which lies a base reality, and that is the ill-typed claim. Counting observers inside models we run tells us how many simulated agents our models contain; it does not bear on whether the axis-free substrate stands beneath our world as a program's base, because no count of models within the world reaches outside it to a substrate of substrates. The arithmetic is about the contents of subsystems, not about the world's ground.

Objection 2 (perhaps there really is a base reality running ours, regress or not). *Response.* "A base reality running ours" is reading (A) again — a computational substrate beneath the world — and is ill-typed for the reasons of Proposition 4.1, not refuted by the present section but diagnosed by the previous one. This section adds only that the *regress* the argument leans on does not exist as a tower of substrates; whether a single base reality could underlie ours is the ill-typed question already addressed. The relocation neither asserts nor denies a base reality; it shows the question to mislocate a projection.

6.3 Section Result and Implications

Nested simulations are nested models in one substrate, not a tower of substrates, so the regress halts and the probabilistic argument inherits the ill-typing of its central claim (Proposition 6.1). The implication is that the hypothesis has no probabilistic engine. We locate the grain of truth it nonetheless feeds on.

7 The Grain of Truth: Representational Isolation

7.1 We meet the world through the model

The simulation hypothesis is compelling, and an account that merely diagnosed it as ill-typed would leave its appeal unexplained. The appeal has a real source, and locating it is the article's last task.

Proposition 7.1 (The grain of truth is representational isolation)**.** *A model-building subsystem meets the world only through its model: its access to the projected world is mediated by the surrogate it maintains, never direct [2, 3]. Hence the world a subsystem experiences is, in a precise sense, already its model — the recovered surrogate, not the substrate unmediated. This representational isolation is the grain of truth the simulation hypothesis feeds on: the sense that we live "inside" a constructed world is correct, but the construction is the receiver-axis model (Article 10), not a computational substrate beneath the world. The truth is about how a subsystem accesses the world, not about what the world rests on.*

Argument. That access is model-mediated is the content of representational isolation, argued in the prior corpus on evolutionary and information-theoretic grounds: a bounded subsystem receiving lossy projections cannot access the world directly and must act on a maintained surrogate [2, 3], the recovering model of Article 11. The world as experienced is therefore the model's, a receiver-axis projection (Article 10), which is exactly the felt situation the simulation hypothesis describes — living within a constructed appearance. The hypothesis errs only in locating the construction beneath the world as a substrate (reading A) rather than within the subsystem as a model (reading B); the construction is real, the relocation corrects its place. □

The proposition explains the hypothesis without conceding it. We do, in a sense, live inside a constructed world: not the substrate met face to face, but the model our subsystem maintains, the only thing a representationally isolated system ever has. That is why the suggestion that reality is "generated," that we inhabit an appearance rather than the thing itself, rings true — it *is* true, on the receiver axis. The simulation hypothesis takes this genuine insight and mislocates it, projecting the model we live inside outward and downward into a computational substrate beneath the world, where it becomes the ill-typed claim of reading (A). The relocation puts the insight back where it belongs: the constructed world we inhabit is our model, in our subsystem, a projection of the one substrate — not a program run by a deeper machine.

7.2 Counterarguments and Responses

Objection 1 (if we only ever access our model, then for all we know the world *is* a simulation). *Response.* Representational isolation means access is model-mediated, not that the world is a simulation in sense (A); the model mediates access to the projected world, a projection of the substrate, not to a deeper computation. "For all we know" invites reading (A) again, and reading (A) is ill-typed, so there is no hidden possibility that isolation leaves open. Isolation is a fact about the receiver axis — we know the world through a model — and it is fully compatible with there being no deeper substrate; indeed the relocation shows the deeper substrate to be not an open possibility but a mislocated projection.

Objection 2 (calling experience "already the model" is idealism, denying the external world). *Response.* It is not idealism: the model is a surrogate *of* the projected world, answerable to the projections the world delivers and corrected against them (Article 11); the external world is presupposed throughout, as the source of the projections the model recovers. Representational isolation concerns the *mediation* of access, not the existence of what is accessed. The world is there, projecting; the subsystem meets it through a model; both are affirmed. The relocation is realist about the substrate and the projected world, and locates only the simulation — the model — in the subsystem.

7.3 Section Result and Implications

The grain of truth in the simulation hypothesis is representational isolation — a subsystem meets the world only through its model, so the world it experiences is already its model — which is a receiver-axis fact, not a deeper substrate (Proposition 7.1). The implication is that the hypothesis's appeal is explained and relocated, not merely dismissed. We record the status of the article's claims.

8 Epistemic Status: Proved, Imported, Affirmed, and What Is Not Claimed

In the manner of the series, we close with an explicit ledger, and we state plainly what the article does not claim. *Proved (here)* marks results established in this article; *Imported* marks results from prior articles; *Affirmed (standard)* marks established work the article locates; *Not claimed* records the positions the article is careful not to take.

Claim	**Status**	**Locus**
Substrate vs projection; reification diagnosis; apophatic boundary	Imported	Kriger, IST 2 [5], Kriger, IST 4 [6]
Recovering model; the simulation remark	Imported	Kriger, IST 11 [10]
"Simulation" has a substrate and a subsystem reading	Proved (here)	Prop. 3.1
World-as-simulation (substrate reading) is ill-typed, not false	Proved (here)	Prop. 4.1

continued on next page

Epistemic-status ledger (continued)

Claim	**Status**	**Locus**
Simulation is the limit of the recovering model, in the subsystem	Proved (here)	Prop. 5.1
Executable map is still a map (type, not output)	Proved (here)	Prop. 5.1
Regress halts: nested models in one substrate	Proved (here)	Prop. 6.1
Grain of truth is representational isolation (receiver axis)	Proved (here)	Prop. 7.1
Bounded systems meet the world only through a model	Imported (prior corpus)	Kriger [2, 3]
The simulation argument (tower, probabilistic inference)	Affirmed (located)	Bostrom [1]
That the world *is* a simulation	Not claimed	—
That the world is *not* a simulation	Not claimed	—

The ledger's last two rows are the article's chief precaution. The relocation does not assert that the world is a simulation, and it does not assert that the world is not a simulation; either assertion would presuppose the ill-typed question well-formed and answer it. The article's claim is the diagnosis: the question, on the substrate reading, mislocates a projection, and the well-typed residue — a subsystem can build a world-sized model that simulates its projection-stream — is true and located in the subsystem. Everything proved here follows the foundations of Articles 2 and 4 and the bridge of Article 11; the prior corpus on representational isolation is imported; and Bostrom's argument is located, neither endorsed nor refuted on its own terms but shown to rest on a mislocation.

9 Discussion

The result of this article is a relocation. The simulation hypothesis, posed as a question about the world — is the world the output of a computation on a deeper substrate? — mislocates a projection. "Simulation" has two readings: a computational substrate beneath the world, and the activity of a model-building subsystem. The first reifies a projection into a second substrate and posits something beneath the axis-free ground, and is therefore ill-typed, non-denoting rather than false. The second is real and is the limit of the recovering model the bridge developed: a model elaborated until running it reproduces the subsystem's projection-stream, an executable map that for all its detail remains a map, located in the subsystem and never the substrate. The simulation, real on the second reading, is relocated from the world to the subsystem, and with it the regress that powered the probabilistic argument dissolves: nested simulations are nested models in one substrate, not a tower of substrates, so there is no stack of grounds to count. The appeal the hypothesis nonetheless exerts is explained by representational isolation, the receiver-axis fact that a subsystem meets the world only through its model, so that the world we experience is already our model — a genuine insight the hypothesis mislocates by projecting it downward into a substrate.

The relocation is of a piece with the volume's method and continues its bridge. Where Article 2 refused to mistake a projection for the substrate and Article 4 marked the predicates that do not cross the block, this article applies both to the most vivid contemporary case, the simulation hypothesis, and finds it an instance of the same mislocation. Where Article 11 established that the model recovers what projection loses, this article follows the recovering model to its limit and finds the simulation there, in the subsystem, exactly where the asymmetry placed it. And where the volume insists throughout that the substrate carries no projection value, this article refuses to grant the world a computational substrate beneath itself, while affirming every well-typed simulation a subsystem can build. The world is not shown to be a simulation, nor shown not to be; the question is shown to be the wrong question, and the right one — what can a model-building subsystem construct, and where does the construction live — is answered: in the subsystem, as a map, however executable.

9.1 Limitations

Three limitations are material. *First*, the article's central verdict is a type-diagnosis, that the world-as-simulation question is ill-typed; a reader who rejects the apophatic boundary of Article 4 — who holds that "a substrate beneath the substrate" is meaningful — will not accept the diagnosis, and the article does not re-argue Article 4 here. *Second*, the relocation affirms that a subsystem can in principle build a world-sized simulation but does not address the practical or thermodynamic limits on doing so; the cost of such a model, bounded by the cost axis (Article 9), is acknowledged but not quantified. *Third*, the grain-of-truth account rests on representational isolation, imported from the prior corpus; its full defense lies in that work, and the article assumes rather than re-establishes it.

9.2 Future Directions

The sequel is the incompleteness article, which takes up the residual floor of Article 11 and the apophatic boundary of Article 4 directly, and to which the present article contributes a case: the ill-typed world-as-simulation question is one of the predicates that do not cross the block, and the incompleteness article can locate it among them. Beyond the volume, the relocation invites a precise question continuous with the receiver axis: whether the richness at which a model becomes a simulation — the point where forward execution reproduces the projection-stream — can be characterized as a threshold on the receiver-relative measures of Article 10, so that "simulation" would acquire a place in the family of receiver-relative projections rather than a merely intuitive sense. That question is left to the work that can address it; the present article's task was the relocation, and the simulation is relocated.

Acknowledgments

The author thanks colleagues of the Information Physics Institute and the Institute of Integrative and Interdisciplinary Research for sustained discussion of the foundations of the series.

References

[1] N. Bostrom. Are you living in a computer simulation? *The Philosophical Quarterly*, 53(211):243–255, 2003.

[2] B. Kriger. Evolutionary and information-theoretic argument for the necessity of representational isolation: why direct perception was never an option for complex systems. IIIR. Zenodo, 2026. doi:10.5281/zenodo.18331202.

[3] B. Kriger. The structural distortion principle: a closed-loop model of perception, attention, and world-maintenance in bounded cognitive systems. IIIR. Zenodo, 2026. doi:10.5281/zenodo.18452700.

[4] B. Kriger. The information substrate: differentiation as the ground of information. IST, Volume VII, Article 1. Zenodo, 2026.

[5] B. Kriger. Substrate and projection: why the universe needs no information. IST, Volume VII, Article 2. Zenodo, 2026.

[6] B. Kriger. The categorial boundary: on predicates that do not cross the block. IST, Volume VII, Article 4. Zenodo, 2026.

[7] B. Kriger. The axis of description: Kolmogorov complexity and the algorithmic projection. IST, Volume VII, Article 6. Zenodo, 2026.

[8] B. Kriger. The axis of cost: Landauer's principle and the infodynamics of DESI DR1. IST, Volume VII, Article 9. Zenodo, 2026.

[9] B. Kriger. The axis of the receiver: Bateson's difference and the pragmatic projection. IST, Volume VII, Article 10. Zenodo, 2026.

[10] B. Kriger. Projection loses, the model recovers: the central asymmetry and the opening of the bridge. IST, Volume VII, Article 11. Zenodo, 2026.

The Incompleteness of Measurement: The Residual Floor as a Theorem of Projective Access

Why a subsystem confined to projections has well-typed structure it can never resolve, and how this differs from the apophatic boundary

Boris Kriger[1,2]

[1] Information Physics Institute
boris.kriger@informationphysicsinstitute.net

[2] Institute of Integrative and Interdisciplinary Research, Toronto
boriskriger@interdisciplinary-institute.org

ORCID: https://orcid.org/0009-0001-0034-2903

The Information Substrate Theory (IST) — Volume VII, Article 14
Part III: The Bridge

Abstract

The opening bridge article showed that a model-building subsystem recovers the projectable structure of the world by inference and leaves a residual floor: a part of the discarded fiber that no accumulation of projections ever resolves, of positive entropy $H_\infty > 0$. This article turns the floor into a theorem. We define a subsystem's projective access as the family of projections available to it, and the joint fiber as the set of substrate configurations that every available projection fails to distinguish; the residual floor is precisely the entropy of the joint fiber. The incompleteness of measurement is then the statement that for any proper projective access — any family of genuine reductions short of the bijective, non-projecting limit — the joint fiber is non-trivial and the floor is positive: a subsystem confined to projections necessarily has structure it cannot measure, not for want of effort or instruments but because projective access has fibers. A coverage construction makes the floor concrete: as projections accumulate the unresolved count falls but plateaus at the distinctions no projection touches, and only the maximal observer that is no subsystem reaches zero. We then draw the distinction the bridge has been preparing. The apophatic boundary of the categorial-boundary article concerns predicates that are ill-typed — questions that do not denote because they read an axis-internal predicate onto the substrate. The incompleteness floor concerns questions that are well-typed — genuine distinctions of the projected world, with determinate answers — yet are forever unresolved by a given access. Two limits bound a subsystem's knowledge from two sides: the categorial limit at the substrate, where questions cease to be well-formed, and the projective limit within the projected world, where well-formed questions cease to be answerable. Finally we note that self-measurement has a fiber too: a subsystem measuring from within is part of what it measures, and complete self-measurement from inside is impossible — a structural cousin, by analogy and not by reduction, of algorithmic incompleteness.

Keywords: incompleteness of measurement; projective access; joint fiber; residual entropy; apophatic boundary; ill-typed predicates; self-measurement; foundations of physics.

Contents

1 Introduction

The opening article of the bridge established the central asymmetry and, with it, a residue. A model-building subsystem recovers the projectable structure of the world by inference rather than by inversion, conditioning an internal surrogate on the projections it receives; but the recovery stops at a floor. There is a part of the discarded fiber that no accumulation of projections ever distinguishes — the part every projection holds in common as a collapsed distinction — and on that part the subsystem's posterior cannot concentrate. The residual uncertainty $H_\infty > 0$ persists as projections accumulate without bound, and it was named, there, the seed of the incompleteness of measurement [6]. This article grows the seed into a theorem.

The construction is direct. A subsystem does not have all projections; it has the ones its organization affords — a definite family, its *projective access*. Each projection in the family is a reduction, collapsing a fiber of substrate configurations it does not distinguish (Article 11); the configurations that *every* projection in the access collapses together form the *joint fiber*, the set of substrate states the subsystem cannot tell apart by any measurement available to it. The residual floor is the entropy of this joint fiber. And the incompleteness of measurement is the statement that, for any proper projective access — any family of genuine reductions that stops short of the bijective, non-projecting limit — the joint fiber is non-trivial, so the floor is positive. A subsystem confined to projections necessarily has structure it cannot measure. The limitation is not contingent, not a matter of finer instruments or longer observation; it is the structural consequence of access being projective, and projections having fibers.

The theorem invites a distinction the bridge has been preparing since Part I, and drawing it is the article's second task. There are two ways a question about the world can fail to be answerable, and they are not the same. The categorial-boundary article identified the first: some predicates are *ill-typed* — they ask of the substrate something only a projection can satisfy, reading an axis-internal predicate (a colour, an age, a position) onto the axis-free ground, and so they do not denote; they are neither true nor false but malformed [4]. The incompleteness floor is the second, and it is different in kind. The questions it concerns are *well-typed*: they ask about genuine distinctions of the projected world, distinctions with determinate answers, configurations that really do or do not obtain. What fails is not their sense but their reachability: no projection in the subsystem's access resolves them. The floor is well-typed and unresolved; the apophatic boundary is ill-typed. One is a limit within the projected world, the other a limit at the substrate.

These two limits bound a subsystem's knowledge from two sides, and seeing them together is the article's point. Above, at the substrate, lies the categorial limit: questions that cease to be well-formed, the apophatic silence of Article 4. Within, in the projected world, lies the projective limit: well-formed questions that cease to be answerable, the incompleteness floor of this article. Neither is a defect to be repaired. The first is the price of the substrate's being axis-free; the second is the price of a subsystem's access being projective. We make the floor a theorem, distinguish it sharply from the apophatic boundary, and close by observing that even a subsystem's measurement of *itself* has a fiber — that complete self-measurement from within is impossible, a structural cousin of algorithmic incompleteness which we draw as an analogy and not a reduction [1].

2 Logical Structure of the Argument

The article moves from the joint fiber to the two boundaries and the self-fiber.

Step 1 (Section 3): Projective access and the joint fiber.
The access is the family of available projections; the joint fiber is what they all collapse; the residual floor is its entropy (Definition 3.1, Proposition 3.2).

Step 2 (Section 4): The incompleteness of measurement.
For proper access the joint fiber is non-trivial and the floor is positive: structure no measurement resolves (Theorem 4.1, Figure 1).

Step 3 (Section 5): Two boundaries.
The incompleteness floor is well-typed-but-unresolved; the apophatic boundary is ill-typed; they are limits of different kinds (Proposition 5.1).

Step 4 (Section 6): Measurement from within.
A subsystem measuring itself is part of what it measures; its self-directed access has a fiber; complete self-measurement is impossible (Proposition 6.1).

Step 5 (Section 7): Epistemic status.
A ledger separates what is proved, imported, and affirmed, and marks the algorithmic analogy as analogy.

The dependencies are linear: Step 1 identifies the floor with the joint fiber; Step 2 proves the fiber non-trivial for proper access; Step 3 distinguishes this projective limit from the categorial one; Step 4 extends the limit to self-measurement. The thesis is that incompleteness of measurement is a theorem of projective access, distinct in kind from the apophatic boundary, and that it reaches even the subsystem's knowledge of itself.

3 Projective Access and the Joint Fiber

3.1 The floor is the joint fiber

We make the residual floor a definite object, the joint fiber of a subsystem's projective access.

Definition 3.1 (Projective access and the joint fiber)**.** A subsystem's *projective access* is the family $\mathcal{P} = \{\rho_i\}_{i \in I}$ of projections available to it. Each ρ_i collapses a fiber of substrate configurations it does not distinguish. The *joint fiber* $\mathcal{F}_\mathcal{P}$ is the set of substrate configurations that *every* ρ_i fails to distinguish: those mapped to the same value by all available projections, $\mathcal{F}_\mathcal{P} = \{x : \rho_i(x) = \rho_i(x') \text{ for all } i \text{ and some } x' \neq x\}$, i.e. the configurations the subsystem cannot tell apart by any measurement it has.

Proposition 3.2 (The residual floor is the entropy of the joint fiber)**.** *The residual floor H_∞ of the recovering model (Article 11) equals the entropy of the joint fiber: $H_\infty = H(X \mid \{\rho_i(X)\}_{i \in I})$, the uncertainty about the substrate configuration X remaining after conditioning on all available projections. For a uniform prior over a finite joint fiber, $H_\infty = \log |\mathcal{F}_\mathcal{P}|$. The model recovers everything the access jointly distinguishes and nothing it jointly collapses.*

Proof. By Article 11 the recovering model conditions on the projections received, and its residual uncertainty about the preimage is the conditional entropy of the substrate configuration given those projections [2, 6]. As projections accumulate within the access, the posterior concentrates on the set of configurations consistent with all of them — the complement of what they jointly distinguish — and cannot fall below the entropy of the configurations they jointly fail to distinguish, which is the joint fiber (Definition 3.1). Hence $H_\infty = H(X \mid \{\rho_i(X)\})$, equal to $\log |\mathcal{F}_\mathcal{P}|$ for a uniform prior. The floor is the joint fiber's entropy. □

The proposition fixes what was, in Article 11, a limit of an accumulating posterior as a definite structural object. The floor is not a vague residue but the joint fiber of the subsystem's access: the equivalence class of substrate configurations that all of the subsystem's projections map alike. Two worlds that differ only within the joint fiber deliver the subsystem identical projections under every measurement it has, and so are, for it, indistinguishable — not provisionally, pending better data, but exactly, because no available projection separates them. The size of that class is the floor.

3.2 Counterarguments and Responses

Objection 1 (a subsystem can acquire new projections, enlarging its access, so the joint fiber is not fixed). *Response.* The joint fiber is relative to an access, and enlarging the access shrinks it: adding a projection that distinguishes some previously collapsed configurations removes them from the joint fiber and lowers the floor (Section 4, Figure 1). The theorem of the next section concerns what remains for *any* proper access, however enlarged; it does not claim a fixed fiber but a positive floor for every access short of the non-projecting limit. The relativity to access is built in and is exactly what the staircase of Figure 1 displays.

Objection 2 (conditional entropy is an information-theoretic quantity, not a fact about what a subsystem can measure). *Response.* The conditional entropy $H(X \mid \{\rho_i(X)\})$ is precisely the uncertainty about X that no function of the available projections can remove, and a measurement available to the subsystem is, by definition, a function of its projections (it has no other access). So the conditional entropy is the exact information-theoretic form of "what the subsystem cannot measure": any estimator it builds is a function of $\{\rho_i(X)\}$, and the data-processing inequality bounds what any such function recovers (Article 11). The quantity and the measurement limit coincide.

3.3 Section Result and Implications

The residual floor is the entropy of the joint fiber of a subsystem's projective access — the configurations all its projections collapse together (Definition 3.1, Proposition 3.2). The implication is that incompleteness, if it holds, is the non-triviality of this joint fiber. We prove it non-trivial.

4 The Incompleteness of Measurement

4.1 The theorem

The floor is the joint fiber's entropy; the theorem is that, for any proper access, this is positive.

Theorem 4.1 (Incompleteness of measurement)**.** *Let $\mathcal{P} = \{\rho_i\}$ be a* proper *projective access: each ρ_i is a genuine reduction (non-trivial fiber), and $\mathcal{P}$ is not the bijective, non-projecting limit that distinguishes every configuration. Then the joint fiber $\mathcal{F}_{\mathcal{P}}$ is non-trivial, $|\mathcal{F}_{\mathcal{P}}| > 1$, and the residual floor is positive, $H_\infty > 0$. A subsystem with proper projective access necessarily has substrate structure it cannot measure — distinctions no available projection resolves — and this is a structural feature of projective access, not a contingent shortfall of instruments or observation.*

Proof. By hypothesis $\mathcal{P}$ is not the bijective limit, so the joint reduction $\rho_{\mathcal{P}} : x \mapsto (\rho_i(x))_{i\in I}$ is not injective: there exist $x \neq x'$ with $\rho_i(x) = \rho_i(x')$ for all i, since otherwise the family would jointly distinguish every pair and $\mathcal{P}$ would be the bijective limit excluded by hypothesis. Hence the joint fiber containing x has at least two elements, $|\mathcal{F}_{\mathcal{P}}| > 1$, and by Proposition 3.2 the floor is $H_\infty = H(X \mid \{\rho_i(X)\}) \geq \log 2 > 0$ for a prior giving positive weight to x and x'. That this holds for *every* proper access establishes the limitation as structural: only the non-projecting limit, which distinguishes all configurations and is therefore no reduction at all, achieves $H_\infty = 0$. □

The theorem says, plainly, that to access the world by projection is to leave some of it unmeasured, and that no enlargement of one's measurements short of becoming the world's bijective image removes the residue. Figure 1 shows the mechanism. Each projection resolves some distinctions and not others; as projections are added the resolved

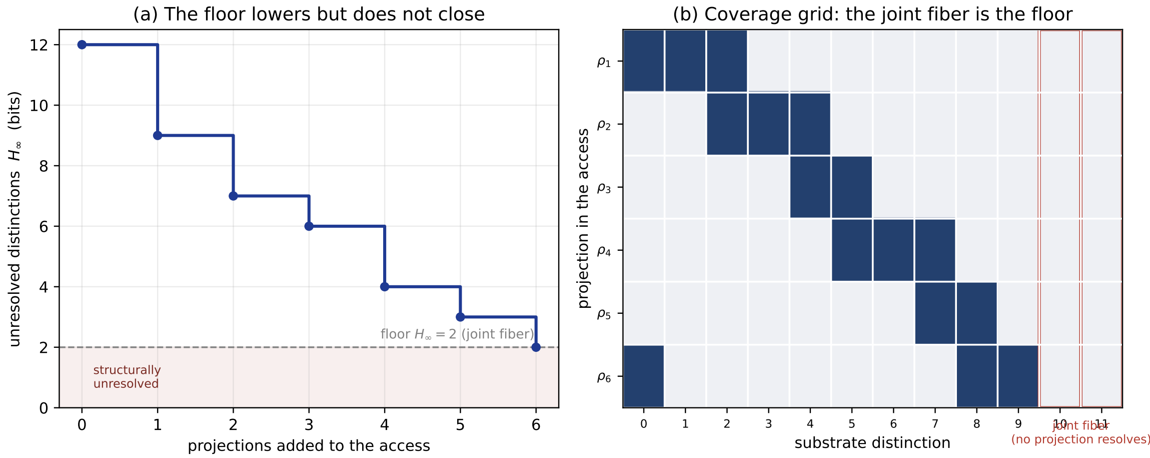

Figure 1: The incompleteness of measurement as the joint fiber. A substrate state carries $D = 12$ independent distinctions; a subsystem's access is a family of projections, each resolving a subset and blind to the rest. **(a)** As projections accumulate the unresolved count H_∞ falls as a staircase but plateaus at the floor — the distinctions no projection touches. More measurement lowers the floor; it does not close it. Only the non-projecting (bijective) limit, the maximal observer that is no subsystem, reaches zero. **(b)** The coverage grid: which distinction (column) each projection (row) resolves. The columns no row touches are the joint fiber (boxed) — the floor, the well-typed structure this access can never resolve.

distinctions accumulate and the unresolved count falls; but the fall plateaus at exactly the distinctions that no projection in the access touches, the joint fiber, and there the floor sits. The only way to drive the floor to zero is to possess a family of projections that jointly distinguish every configuration — the bijective limit, which is not a projective access at all but the maximal observer of Article 10, the receiver that reduces nothing. A subsystem, which is by definition a reducing receiver, never occupies that limit, and so always has a positive floor.

4.2 Counterarguments and Responses

Objection 1 (with enough independent projections the floor can be made as small as one likes, so "incompleteness" overstates a quantitative smallness). *Response.* The floor can indeed be made small by a rich access, and the theorem asserts not that it is large but that it is positive for every proper access — that it never reaches zero. The claim is qualitative and exact: there is always some well-typed structure unresolved, however much is resolved. "Incompleteness" names this strict positivity, the impossibility of total measurement, not a claim about magnitude; a subsystem may know almost everything and still, necessarily, not everything.

Objection 2 (the bijective limit is achievable in principle, so the floor is not necessary). *Response.* The bijective limit distinguishes every configuration and so collapses no fiber; it is the non-reducing receiver, which Article 10 identified as the maximal observer and Article 11 noted is not available to a subsystem inside the world, since a subsystem's access is finite and organized and therefore reduces. A receiver that reduced nothing would have to register every distinction of the entire substrate, which is not a subsystem's situation but the limit in which receiver and world coincide. For any actual subsystem the limit is unattained, and the floor is positive — which is the theorem.

4.3 Section Result and Implications

For any proper projective access the joint fiber is non-trivial and the floor is positive: a subsystem confined to projections necessarily has well-typed structure it cannot measure (Theorem 4.1, Figure 1). The implication is that this projective limit must be distinguished from the categorial limit of Article 4. We distinguish them.

5 Two Boundaries: Apophatic and Projective

5.1 Ill-typed and well-typed-but-unresolved

The incompleteness floor is a limit on knowledge, and so is the apophatic boundary of Article 4; the article's second task is to keep them apart, for they are different in kind.

Proposition 5.1 (Two boundaries of different kinds)**.** *A subsystem's knowledge is bounded in two distinct ways:*

(i) (Categorial / apophatic.) *The apophatic boundary concerns* ill-typed *questions: predicates applied to the axis-free substrate that only a projection can satisfy. These do not denote; they are neither true nor false but malformed (Article 4). The limit is at the substrate.*

(ii) (Projective / incompleteness.) *The incompleteness floor concerns* well-typed *questions: genuine distinctions of the projected world, with determinate answers, that no projection in a given access resolves (Theorem 4.1). These denote and have answers; the answers are unreachable. The limit is within the projected world.*

The two are not the same boundary at different strengths; one marks where questions cease to be well-formed, the other where well-formed questions cease to be answerable.

Proof. The apophatic boundary is the set of predicates whose application to the substrate is a category error (Article 4): to ask the substrate's colour or age is to apply an axis-internal predicate to the axis-free ground, which fails to denote. The incompleteness floor, by contrast, consists of distinctions *within* the projected world — configurations x, x' that genuinely differ (Theorem 4.1) and so pose a well-formed question, "is it x or x'?", with a determinate answer fixed by which obtains. That question denotes and has a truth-value; what the joint fiber establishes is only that no available projection delivers it. The first failure is of sense, the second of access; a malformed non-question and a well-formed unanswerable question are different, so the boundaries differ in kind. □

The distinction matters because the two limits are easily conflated, and conflating them distorts both. Treat the incompleteness floor as apophatic, and one wrongly concludes that the unmeasured distinctions are not real — that because we cannot resolve them there is no fact of the matter — when the theorem says they are perfectly real distinctions of the projected world, merely beyond a given access. Treat the apophatic boundary as incompleteness, and one wrongly imagines that the substrate's "colour" is a hidden fact awaiting a better instrument, when Article 4 showed there is no such fact to find, the question being malformed. The bridge keeps them separate: the substrate is apophatically silent, not incompletely measured; the projected world is incompletely measured, not apophatically silent. The two silences have different sources and different honest responses — to the first, that the question does not denote; to the second, that the answer exists and is unreachable.

5.2 Counterarguments and Responses

Objection 1 (an unresolvable distinction is operationally meaningless, so the floor collapses into the apophatic after all). *Response.* Operationally inaccessible is not meaningless. The configurations x and x' in the joint fiber differ as substrate configurations and would be distinguished by some projection outside the access — indeed by the bijective limit — so the distinction is well-defined and the question "x or x'?" is well-formed, with an answer. A verificationist who equates meaning with the subsystem's verifiability will indeed collapse the two; the substrate frame does not, holding that what a particular access cannot resolve may still be a genuine distinction of the world. The floor's distinctions are real and unreached, which is exactly not the apophatic case of no-fact-to-find.

Objection 2 (if the floor's distinctions can never be resolved, asserting they are "real" is idle metaphysics). *Response.* It is not idle: it is the difference between two diagnoses with different consequences. That the floor's distinctions are real and unreached means a richer access *would* resolve them (Figure 1), so the limit is relative to access and can be lowered, whereas the apophatic limit cannot be lowered by any access because there is nothing to resolve. The reality of the floor's distinctions is what makes the floor a function of access rather than an absolute silence, and that is a substantive, not idle, difference — it is why measurement can progress at all.

5.3 Section Result and Implications

The incompleteness floor (well-typed, unresolved, within the projected world) and the apophatic boundary (ill-typed, non-denoting, at the substrate) are limits of different kinds, bounding a subsystem's knowledge from two sides (Proposition 5.1). The implication is that the projective limit reaches even inward. We turn it on the subsystem itself.

6 Measurement from Within

6.1 The subsystem's self-fiber

The incompleteness floor concerns a subsystem's access to the world; the subsystem is part of the world, so the floor reaches its access to itself.

Proposition 6.1 (Self-measurement has a fiber)**.** *A subsystem that measures from within is part of what it measures: its self-directed projections are taken by a part of the system about the whole, including the measuring part. Such self-access is a proper projective access (it reduces, being finite and organized), so by Theorem 4.1 its joint fiber is non-trivial: there is structure of the subsystem itself — including of its own measuring apparatus — that the subsystem cannot resolve from within. Complete self-measurement from inside is therefore impossible. This is a structural cousin, by analogy and not by reduction, of algorithmic incompleteness, in which a sufficiently rich formal system cannot certify all truths about itself [1].*

Argument. A subsystem's self-measurement is a projection whose domain includes the subsystem, taken by a proper part of it; the apparatus that records cannot, in the act, fully project its own recording state, since registering a state changes it and the registration of the registration regresses. The self-access is thus a proper projective access (it omits at least the current state of its own apparatus), and Theorem 4.1 gives a non-trivial joint fiber: self-structure the subsystem cannot resolve from within. The analogy to algorithmic incompleteness is that a system rich enough to model itself cannot, from inside, settle every question about itself; we mark it as an analogy, the present result resting on the projective fiber of self-access, not on a formal arithmetization [1]. □

The proposition extends the limit to its most reflexive case and tempers a hope. One might suppose that whatever a subsystem cannot measure of the external world it could at least measure of itself, having privileged inner access; the proposition denies it. Self-access is projective too — a part registering the whole — and so has its own joint fiber, its own floor. The measuring apparatus cannot fully measure its own measuring state in the act of measuring, and the regress of measuring the measurement does not close. A subsystem is incompletely transparent to itself for the same structural reason it is incompletely informed about the world: its access, inward or outward, is projective, and projections have fibers. We draw the kinship with algorithmic incompleteness carefully, as an analogy of structure — a system rich enough to represent itself meets a limit on self-certification — and not as a derivation of one from the other; the present limit is information-theoretic and projective, Chaitin's is computational and arithmetical, and the resemblance is in shape, not in proof.

6.2 Counterarguments and Responses

Objection 1 (a subsystem could be measured completely by an external observer, so the limit is not absolute). *Response.* An external observer with a richer access could resolve distinctions the subsystem cannot resolve of itself, exactly as a richer access lowers any floor (Figure 1); the proposition concerns self-measurement *from within*, not measurement in principle. The limit is on the subsystem's own access to itself, which is proper and so has a fiber; it is not a claim that the subsystem's states are absolutely unmeasurable, only that they are not fully self-measurable. The relativity to access holds inward as outward.

Objection 2 (the algorithmic-incompleteness analogy is doing illegitimate work). *Response.* It is doing no probative work and is flagged as analogy in the statement, the proof, and the ledger. The proposition rests entirely on the projective fiber of self-access (Theorem 4.1); the mention of algorithmic incompleteness only notes a structural resemblance — a self-modelling system limited in self-certification — to orient the reader, and explicitly disclaims a reduction. Removing the analogy would not weaken the proof, which is information-theoretic throughout.

6.3 Section Result and Implications

Self-measurement is projective and so has a joint fiber: complete self-measurement from within is impossible, a structural cousin of algorithmic incompleteness drawn as analogy (Proposition 6.1). The implication is that the projective limit is fully general, reaching the world and the subsystem alike. We record status and close.

7 Epistemic Status: Proved, Imported, Affirmed

In the manner of the series, we close with an explicit ledger. *Proved (here)* marks results established in this article; *Imported* marks results from prior articles; *Affirmed (standard)* marks established work the article applies; *Analogy (not reduction)* marks the algorithmic kinship.

Claim	Status	Locus
Residual floor $H_\infty > 0$ of the recovering model	Imported	Kriger, IST 11 [6]
Apophatic boundary; ill-typed predicates	Imported	Kriger, IST 4 [4]
Maximal observer as non-reducing receiver	Imported	Kriger, IST 10 [5]
Projective access; joint fiber	Proved (here)	Def. 3.1
Floor equals entropy of the joint fiber	Proved (here)	Prop. 3.2
Incompleteness: proper access $\Rightarrow H_\infty > 0$	Proved (here)	Thm. 4.1
Floor lowers with access but never closes	Proved (here)	Fig. 1
Two boundaries: well-typed-unresolved vs ill-typed	Proved (here)	Prop. 5.1
Self-measurement from within is incomplete	Proved (here)	Prop. 6.1
Conditional entropy; data-processing bound	Affirmed (standard)	Cover and Thomas [2]
Bounded systems access the world projectively	Imported (prior corpus)	Kriger [3]
Kinship with algorithmic incompleteness	Analogy (not reduction)	Chaitin [1]

The ledger marks the division of labor. The article *imports* the residual floor, the apophatic boundary, and the maximal observer; it *proves* that the floor is the joint fiber, that proper access entails a positive floor, that the floor lowers but never closes, that the projective limit differs in kind from the apophatic, and that self-measurement is incomplete; it *affirms* the standard information theory it applies; and it marks the algorithmic-incompleteness kinship explicitly as an analogy, not a reduction.

8 Discussion

The result of this article is a theorem and a distinction. The theorem turns the residual floor of the opening bridge article into a structural fact: a subsystem's projective access has a joint fiber, the configurations all its projections collapse together, and for any proper access this fiber is non-trivial, so the floor H_∞ is positive. A subsystem confined to projections necessarily has well-typed structure it cannot measure, and no enlargement of its measurements short of the non-projecting limit — the maximal observer that is no subsystem — removes the residue. The distinction sets this projective limit beside the categorial one of Article 4 and keeps them apart: the apophatic boundary concerns ill-typed questions that do not denote, a limit at the substrate; the incompleteness floor concerns well-typed questions that denote and have answers but are unreachable, a limit within the projected world. Two boundaries, of different kinds, bound a subsystem's knowledge from two sides — the silence of malformed questions above, the silence of unanswerable well-formed questions within.

The two limits complete the bridge's account of the model-builder. Article 11 showed that projection loses and the model recovers by inference; Article 12 found the arrow of time in the model-builder's record gradient; Article 13 relocated the simulation into the model; and this article marks the outer edge of what the model can recover at all. The recovering model is powerful — it climbs the entropy gradient laying down records, it can become a simulation rich enough to reproduce its projection-stream — but it is bounded, above by the apophatic silence of the substrate and within by the incompleteness floor of its own access. The model-builder is neither omniscient nor deluded: it recovers the projectable structure of its world, truly and progressively, and it meets, at the joint fiber, a limit that is not a failure but the shape of projective access itself. That even its self-measurement is incomplete — that it cannot be fully transparent to itself from within — places the limit at the most intimate possible point and confirms its generality. The prior corpus had argued that bounded systems access the world only projectively, through a maintained model [3]; the present article draws the consequence: projective access, inward or outward, leaves a floor, and the floor is where measurement ends.

8.1 Limitations

Three limitations are material. *First*, the theorem establishes that the floor is positive for proper access but does not, in general, compute its size; the magnitude H_∞ depends on the joint fiber of a particular access and is not given by the theorem, which asserts only strict positivity. *Second*, the distinction between the apophatic and projective boundaries presupposes the categorial boundary of Article 4; a reader who rejects that boundary will read both limits as one, and the article does not re-argue Article 4 here. *Third*, the self-measurement result is argued at the level of the projective fiber of self-access and drawn as an analogy to algorithmic incompleteness; a precise formal correspondence, if one exists, is not established and is not claimed.

8.2 Future Directions

The sequel is the volume's closing treatment of the apophatic boundary, which can now incorporate both limits — the categorial silence of the substrate and the projective floor of access — into a single account of what lies beyond a subsystem's reach and why. Beyond the volume, the theorem invites the precise question the series has raised at each stage: whether the floor H_∞ of an access is a projection-theoretic invariant, computable from the joint fiber's relation to the inter-axis morphisms of Article 3, so that the limit of what any model can recover would be read off the same diagram that classifies the axes. The completeness conjecture and the incompleteness floor would then be two faces of one structure — what the morphisms preserve and what the joint fiber discards. That question is left to the work that can address it; the present article's task was to make the floor a theorem, and the floor is a theorem.

Acknowledgments

The author thanks colleagues of the Information Physics Institute and the Institute of Integrative and Interdisciplinary Research for sustained discussion of the foundations of the series.

References

[1] G. J. Chaitin. A theory of program size formally identical to Gödel's incompleteness theorem. *Journal of the ACM*, 22(3):329–340, 1975.

[2] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Wiley-Interscience, Hoboken, NJ, 2nd edition, 2006.

[3] B. Kriger. Evolutionary and information-theoretic argument for the necessity of representational isolation: why direct perception was never an option for complex systems. IIIR. Zenodo, 2026. doi:10.5281/zenodo.18331202.

[4] B. Kriger. The categorial boundary: on predicates that do not cross the block. IST, Volume VII, Article 4. Zenodo, 2026.

[5] B. Kriger. The axis of the receiver: Bateson's difference and the pragmatic projection. IST, Volume VII, Article 10. Zenodo, 2026.

[6] B. Kriger. Projection loses, the model recovers: the central asymmetry and the opening of the bridge. IST, Volume VII, Article 11. Zenodo, 2026.

A Code for Figure 1

The following Python script generates Figure 1: a substrate carrying $D = 12$ distinctions, a projective access of six projections, and the joint fiber — the distinctions no projection resolves — as the residual floor.

```
D = 12
access = [{0,1,2}, {2,3,4}, {4,5}, {5,6,7}, {7,8}, {8,9,0}]   # each resolves a subset
# distinctions 10, 11 are resolved by NO projection -> the joint fiber (floor = 2)

covered, unresolved = set(), [D]
for S in access:
    covered |= S
    unresolved.append(D - len(covered))      # staircase: 12, 9, 7, 6, 4, 3, 2

joint_fiber = [d for d in range(D) if all(d not in S for S in access)]  # [10, 11]
floor = len(joint_fiber)                                                # H_inf = 2
```

The Two Silences:

The Apophatic and Projective Boundaries of a Subsystem's Reach

One boundary at the substrate where questions cease to be well-formed, one within the world where well-formed questions cease to be answerable

Boris Kriger[1,2]

[1] Information Physics Institute
boris.kriger@informationphysicsinstitute.net

[2] Institute of Integrative and Interdisciplinary Research, Toronto
boriskriger@interdisciplinary-institute.org

ORCID: https://orcid.org/0009-0001-0034-2903

The Information Substrate Theory (IST) — Volume VII, Article 15
Part III: The Bridge — toward the close

Abstract

The volume has found two limits on what a subsystem can know, and this article, beginning the close, draws them into a single account. The first was established at the foundations: the categorial boundary, where predicates applied to the axis-free substrate fail to denote, so that certain questions are not unanswered but ill-formed. The second was established as a theorem of the bridge: the projective boundary, where well-formed questions about genuine distinctions of the projected world go forever unresolved because a subsystem's access is projective and projections have fibers. The two are silences of different kinds, and the article's chief work is to keep them distinct while exhibiting them as the single boundary of a subsystem's reach. We give the map: between the two silences lies the rich interior of the recovered and projectable, the whole content of Parts II and III; the boundaries bound that interior, they do not empty it. We then state the discipline of speech proper to each. At the categorial boundary only the via negativa is well-formed: of the substrate one may say what it is not, and what it stands in relation to, but never a positive axis-value, since every such filling is the reification the foundations forbid. At the projective boundary a different discipline holds: one asserts that there is a determinate fact and that it is unreached by the present access, lowerable by a richer access but never closable short of the non-projecting limit; one must not conclude that there is no fact, nor that some instrument will surely reach it. The two errors to avoid are symmetric — treating the substrate as a hidden fact awaiting a better measurement, and treating the unreached distinctions as no facts at all — and both collapse the two silences into one. The account is neither skeptical nor mystical: it is a disciplined map of a vast knowable interior and the two different edges that bound it. The kinship of the categorial silence to the apophatic method of older traditions is noted and, as firmly, distinguished: the substrate's ineffability is its axis-freeness, fully explicable, not a mystery.

Keywords: apophatic boundary; projective boundary; via negativa; ill-typed predicates; incompleteness; limits of knowledge; substrate; foundations of physics.

Contents

1 Introduction

The volume has arrived, by two routes, at two limits, and the work of this article is to set them side by side. The first limit was fixed at the foundations. The categorial-boundary article showed that predicates belonging to an axis — a colour, an age, a position, a measure — cannot be applied to the axis-free substrate without category error; questions that so apply them do not denote, and are therefore not unanswered but ill-formed, neither true nor false but malformed [4]. The second limit was proved on the bridge. The incompleteness article showed that a subsystem's projective access has a joint fiber — configurations every available projection collapses together — and that for any proper access this fiber is non-trivial, so there is well-typed structure, genuine distinctions of the projected world with determinate answers, that no measurement the subsystem has ever resolves [7]. Two limits, two silences, reached from opposite directions: one at the edge of sense, one at the edge of reach.

They are easy to confuse and important not to, and the confusion runs both ways. One may take the categorial silence for an incompleteness — imagining that the substrate's "colour" or "age" is a hidden fact awaiting a finer instrument — when the categorial article showed there is no such fact to find, the question being malformed. Or one may take the projective floor for an apophatic silence — concluding that because some distinction can never be measured there is no fact of the matter — when the incompleteness theorem showed the distinction is perfectly real, a genuine difference of the projected world, merely beyond a given access. The two errors are symmetric, and each collapses the two silences into one. Keeping them apart, and saying exactly what discipline of speech each demands, is this article's task.

The map is simple and worth stating at once. Between the two silences lies the interior: the recovered and projectable, everything Part II classified as the projections and Part III showed the model-builder to recover, time's arrow and the relocated simulation included. This interior is vast and real; the boundaries bound it, they do not empty it. At its outer edge stands the categorial boundary, beyond which is the axis-free substrate, approached only by the via negativa — one may say what the substrate is *not*, and what it stands in relation to, but never a positive axis-value, since every positive filling reifies a projection into the substrate, the error the foundations forbid [2, 3]. Within the interior, between what a subsystem has reached and what its world contains, runs the projective boundary, the joint fiber of its access: well-typed distinctions, real and determinate, that this access never resolves, lowerable by a richer access but never closable short of the non-projecting limit that is no subsystem [5–7].

Two boundaries, then, of different kinds and different disciplines, bounding one rich interior. The article gives the map (Section 3), states the via negativa proper to the categorial edge (Section 4) and the real-and-unreached discipline proper to the projective edge (Section 5), names the symmetric double error that collapses them (Section 6), and insists that the result is neither skeptical nor mystical but a disciplined account of a knowable world and its two honest limits (Section 7). The kinship of the categorial silence to the apophatic method of older traditions of thought is real and is noted; but it is, as firmly, distinguished, for the substrate's ineffability is nothing mysterious — it is simply its axis-freeness, fully explicable, the precise reason no axis-predicate applies.

2 Logical Structure of the Argument

The article unifies the two limits, states the discipline of each, and guards against two errors.

Step 1 (Section 3): The boundary has two components.
The categorial and projective limits are the two edges of a subsystem's reach, bounding a rich interior (Proposition 3.1, Figure 1).

Step 2 (Section 4): The via negativa.
At the categorial boundary only negative and relational predication of the substrate is well-formed; positive axis-predication is reification (Proposition 4.1).

Step 3 (Section 5): Real and unreached.
At the projective boundary one asserts a determinate, unreached, lowerable-but-not-closable fact (Proposition 5.1).

Step 4 (Section 6): The double error.
Collapsing either silence into the other is an error; the two differ in source and discipline (Proposition 6.1).

Step 5 (Section 7): The interior is rich.
The boundaries bound a vast knowable interior; the account is neither skeptical nor mystical (Proposition 7.1).

Step 6 (Section 8): Epistemic status.
A ledger separates what is proved, imported, and affirmed.

The dependencies are linear: Step 1 sets the two edges side by side; Steps 2 and 3 give each its discipline; Step 4 forbids their conflation; Step 5 affirms the interior they bound. The thesis is that a subsystem's reach has two boundaries of different kinds, each demanding its own discipline of speech, and that together they bound a rich knowable world rather than emptying it.

3 The Boundary Has Two Components

3.1 Two edges of one reach

We first set the two limits, separately established, into one structure.

Proposition 3.1 (The boundary of a subsystem's reach has two components)**.** *The boundary of what a subsystem can know comprises two edges of different kinds:*

(i) *the* categorial *(apophatic) boundary, at the axis-free substrate, where predicates fail to denote and questions are ill-formed (Article 4); and*

(ii) *the* projective *boundary, within the projected world, where well-typed questions about real distinctions go unresolved by a given projective access (Article 14).*

Between them lies the interior*: the recovered and projectable, the content of Parts II and III. The two edges bound this interior; they do not constitute its whole, and they do not empty it.*

Proof. The categorial boundary is established in Article 4 as the set of ill-typed predications of the substrate [4]; the projective boundary in Article 14 as the joint fiber of a proper access, of positive residual entropy [7]. They are of different kinds: one concerns sense (whether a question denotes), the other reach (whether a well-formed question is answerable), as Article 14 already distinguished. The interior is what lies inside both: the projections classified in Part II and the structure the model-builder recovers in Part III, all well-typed and all reached. Its non-emptiness is the content of those parts. Hence the boundary of reach has exactly these two components, with the interior between. □

The map (Figure 1) shows the structure at a glance and corrects a reflex. The reflex, when one hears of limits on knowledge, is to picture a small lit clearing in a vast darkness — knowledge as a precarious exception. The map inverts the picture: the lit interior is vast, comprising every projection and every recovery the volume has treated, and the darkness is two thin edges of different kinds, one where sense gives out and one where reach gives out. The boundaries are real and the article will be precise about each, but they are edges of a rich country, not the rumor of one. To mistake the edges for the whole — to read the volume as a counsel of ignorance — is the first thing the map forbids.

3.2 Counterarguments and Responses

Objection 1 (two boundaries are one too many; a single notion of "the unknowable" would do). *Response.* A single notion blurs exactly what matters. The unknowable-because-malformed and the unknowable-because-unreached call for opposite responses: to the first, that there is nothing to know, the question being ill-formed; to the second, that there is something to know, determinate and real, which this access cannot reach. Lumping them as "the unknowable" licenses the two errors of Section 6, treating real distinctions as non-facts or malformed questions as hidden facts. Two boundaries are the minimum the distinctions require.

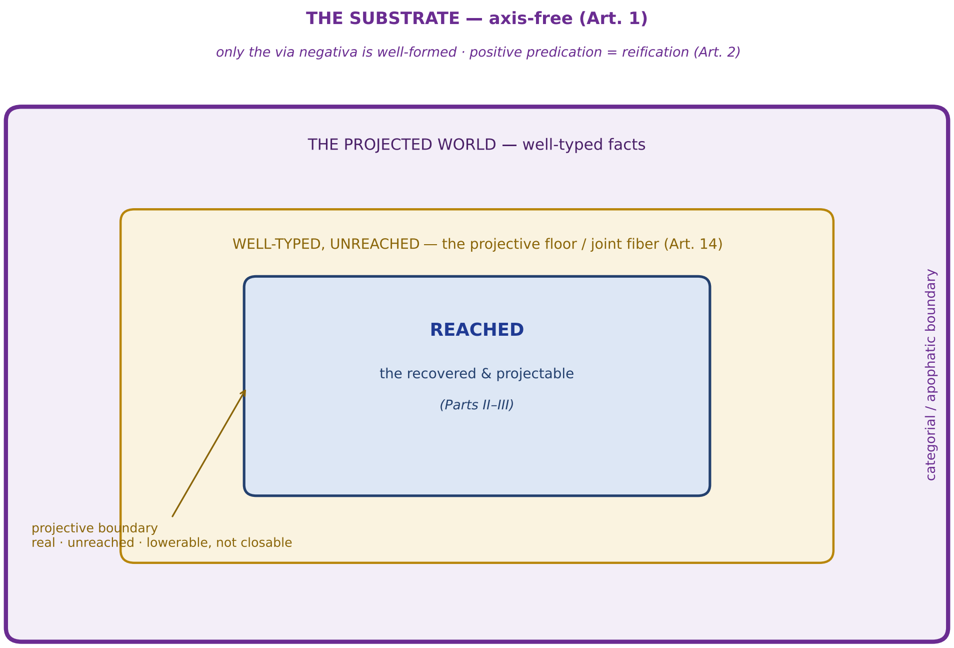

Figure 1: The map of a subsystem's reach. The innermost region is REACHED: the recovered and projectable, the content of Parts II–III. The surrounding ring is WELL-TYPED, UNREACHED: the projective floor, the joint fiber of the subsystem's access (Article 14) — real distinctions of the projected world this access never resolves. The inner edge between them is the *projective boundary*: real, unreached, lowerable by a richer access but not closable. The outer frame is the projected world of well-typed facts; its border is the *categorial / apophatic boundary* (Article 4). Outside everything is the axis-free substrate (Article 1), of which only the via negativa is well-formed and positive predication is reification (Article 2). Two boundaries of different kinds bound a rich interior, not an empty one.

Objection 2 (the interior may be less rich than claimed; perhaps most of the world is in the floor). *Response.* The size of the floor relative to the interior is a quantitative matter the incompleteness theorem leaves open (Article 14), and the article does not claim the interior is most of the world by measure. It claims the interior is rich in the sense that matters — it contains every projection and recovery the volume treats, the entire positive content of physics as the volume reads it — and that the boundaries bound rather than constitute it. A large floor would mean much is real and unreached, which is a claim about reach, not a demotion of the interior's reality or richness.

3.3 Section Result and Implications

A subsystem's reach is bounded by two edges of different kinds, the categorial and the projective, with a rich interior between (Proposition 3.1, Figure 1). The implication is that each edge demands its own discipline of speech. We give them.

4 The Discipline at the Categorial Boundary: the Via Negativa

4.1 What may be said of the substrate

At the outer edge, where predicates meet the axis-free substrate, the discipline is the via negativa.

Proposition 4.1 (Only the via negativa is well-formed of the substrate)**.** *Of the substrate, the well-formed predications are exactly the negative and the relational. One may say truly what the substrate is* not *— not in time, not in space, not coloured, not a measure, not a code — for each such denial removes an axis-predicate that does not apply (Article 4). One may say truly what the substrate stands in* relation *to — that it is that of which the axes are projections, that differentiatedness is its character (Article 1). One may not say what the substrate positively* is *in any axis-value, for every such positive predication applies an axis-internal predicate to the axis-free ground and so reifies a projection into the substrate (Article 2). The mode proper to the categorial boundary is therefore apophatic: approach by negation and relation, never by positive axis-predication.*

Proof. A positive axis-predication of the substrate — "the substrate is [an axis-value]" — applies a predicate that has sense only within an axis to the ground that is prior to every axis, which Article 4 marks ill-typed and Article 2 diagnoses as reification [3, 4]. A negative predication — "the substrate is not [an axis-value]" — is well-formed and true precisely because the axis-predicate does not apply: to deny an inapplicable predicate is correct, not malformed. A relational predication — "the substrate is that of which axes are projections" — is well-formed because it speaks of the substrate's relation to the axes, not of its possessing an axis-value, and is the substrate's positive characterization established in Article 1 [2]. Hence the well-formed modes are the negative and the relational, and the positive-axial mode is excluded; the discipline is the via negativa. □

The proposition gives the categorial boundary its rule of speech, and the rule is exacting. The temptation at this boundary is to fill the silence with a positive picture — to say the substrate *is* a field, or a computation, or information, or a mind, or pure potential — and the temptation is strong because the human demand for a positive answer is strong. Every such filling is a reification: it takes a projection (the field, the computation, the information) and installs it as the substrate, exactly the error Article 2 diagnosed and the relocation of the simulation (Article 13) found at work in the simulation hypothesis. The via negativa is the discipline that resists the temptation. It does not leave the substrate uncharacterized — the relational mode says the true and substantive thing, that the substrate is the differentiated ground of which the axes are projections — but it refuses the positive axis-value the demand keeps asking for, because there is none to give: not a missing one, but a mis-typed one.

4.2 The kinship to the apophatic tradition, and its limit

The term is borrowed deliberately, and the borrowing should be neither hidden nor overread. Older traditions of thought developed an apophatic or negative method for approaching what they took to be ultimate — speaking of it only by negation, refusing positive predication, holding that the ground exceeds every category. The structure of the categorial boundary is kin to this method, and naming the kinship is honest. But the kinship is structural only, and the difference is decisive. In those traditions the refusal of positive predication marked a mystery — the ground was held to exceed the understanding, ineffable in principle. Here nothing is mysterious. The substrate's resistance to positive axis-predication is not an excess beyond understanding but a precise, fully explicable fact: it is axis-free, and axis-predicates apply only within axes, so they do not apply to it — as a matter of type, not of depth. The "ineffability" of the substrate is its axis-freeness, stated plainly; there is no remainder of mystery once the type is understood. The via negativa here is a logical discipline, not a mystical one, and the article claims the structure while declining the mystery.

4.3 Counterarguments and Responses

Objection 1 (the relational characterization "ground of the axes" is itself a positive predication, so the via negativa is not consistently kept). *Response.* It is positive but not axial: it predicates of the substrate a relation to the axes, not possession of an axis-value, and that is exactly the line the proposition draws. "Is the ground of which axes are projections" ascribes no colour, age, position, or measure; it locates the substrate relative to the projections, which is well-formed (Article 1). The via negativa excludes positive *axis-values*, not all positive speech; relational characterization is the permitted positive mode, and it is consistent with the discipline.

Objection 2 (if no positive axis-value applies, the substrate is empty, a mere word). *Response.* It is not empty: it is differentiated, the ground of all the distinctions the axes project (Article 1), and this is a substantive character, not a word. What it lacks is an axis-value, because axis-values are projections and the substrate is what projects; lacking a projection-value is not being empty but being prior to projection. The reflex that equates "no positive axis-value" with "nothing" is the reification reflex again, demanding the substrate be some projection or none; the substrate is neither a projection nor nothing, but the differentiated ground of projections.

4.4 Section Result and Implications

At the categorial boundary the well-formed modes are the negative and the relational; positive axis-predication is reification, and the discipline is the via negativa, a logical and not a mystical one (Proposition 4.1). The implication is that the inner boundary demands a different discipline, since there the questions denote. We give it.

5 The Discipline at the Projective Boundary: Real and Unreached

5.1 What may be said of the floor

At the inner edge, where well-formed questions meet the joint fiber, the discipline is the opposite of silence.

Proposition 5.1 (At the projective boundary: real, unreached, lowerable not closable)**.** *Of a distinction in the projective floor — a difference no projection in the present access resolves (Article 14) — the well-formed assertions are three, held together: that there* is *a determinate fact of the matter (the configurations genuinely differ); that it is* unreached *by the present access (no available projection resolves it); and that it is* lowerable but not closable *— a richer access would resolve it, yet no proper access closes the floor entirely. The two assertions to refuse are symmetric: that there is* no fact *(verificationism, denying the real distinction), and that some instrument will* surely reach *it (denying the floor's persistence under proper access).*

Proof. That there is a determinate fact is the incompleteness theorem of Article 14: the floor's configurations genuinely differ, posing a well-formed question with an answer [7]. That it is unreached is the definition of the joint fiber: no available projection resolves it. That it is lowerable but not closable is the staircase of Article 14: a richer access removes some floor distinctions, but for any proper access the joint fiber remains non-trivial [7]. The refused assertions contradict these: "no fact" denies the real distinction the theorem establishes; "surely reachable" denies the strict positivity of the floor for proper access. Holding the three and refusing the two is the discipline. □

The proposition gives the projective boundary a discipline as exacting as the via negativa but pointed the other way. Here one must *not* fall silent, for the questions denote and have answers; the honest acknowledgment is that the answer exists and this access does not reach it. The verificationist temptation — to say that what cannot be measured is no fact — is the error to resist, and it is the mirror of the reification temptation at the outer boundary: where reification fills the apophatic silence with a false positive, verificationism empties the projective fact into a false silence. Equally to be resisted is the opposite optimism, that any unreached fact is surely reachable with better instruments; the incompleteness theorem forbids it for proper access, since the floor lowers but never closes. The disciplined stance is the threefold one: real, unreached, lowerable-not-closable — neither denying the fact nor promising its capture.

5.2 Counterarguments and Responses

Objection 1 ("lowerable but not closable" is unstable: if every floor distinction is reachable by *some* richer access, then none is permanently unreached). *Response.* Each floor distinction is reachable by some richer access, yet no single proper access reaches all of them, and a subsystem has one access at a time, enlarged only by becoming a different, still proper, subsystem. The distinction reached by enlarging the access is removed from the floor, but the enlarged access has its own joint fiber, its own floor (Article 14). So "lowerable but not closable" is stable: any given distinction may be reached, but the floor as such is never emptied, because every proper access has one. The relativity to access is the whole point, not an instability.

Objection 2 (asserting a fact one can never reach is idle, indistinguishable in practice from denying it). *Response.* It is not idle, and the practical difference is real: the assertion that the fact is real and unreached licenses the search for a richer access that would reach it (the floor is lowerable), whereas denying the fact forecloses the search. A century of physics has lowered floors that earlier looked absolute by finding new projections; that progress presupposes the disciplined stance, that the unreached is real and possibly-reachable, not the verificationist one that the unmeasured is unreal. The stance is what makes inquiry rational rather than complacent.

5.3 Section Result and Implications

At the projective boundary the discipline is to assert a real, unreached, lowerable-not-closable fact, refusing both verificationist denial and optimistic capture (Proposition 5.1). The implication is that the two boundaries demand opposite disciplines, and conflating them confuses the disciplines. We name the error.

6 The Double Error: Why the Two Must Not Be Collapsed

6.1 The symmetric mistake

With each boundary's discipline in hand, the error of collapsing them can be stated exactly, and it is symmetric.

Proposition 6.1 (The double error)**.** *Collapsing the two silences into one produces a symmetric pair of errors:*

(i) (Categorial as projective.) *Treating the categorial boundary as an incompleteness — supposing the substrate has a positive axis-value that a finer measurement would reach — imports the projective discipline where the apophatic belongs, and so reifies (Article 2): it seeks a hidden fact where there is none, the question being ill-formed.*

(ii) (Projective as categorial.) *Treating the projective floor as apophatic — supposing the unreached distinctions are no facts at all — imports the apophatic discipline where the projective belongs, and so falls into verificationism: it denies a real fact because the present access cannot reach it.*

The two boundaries differ in source — the axis-freeness of the substrate versus the projectivity of access — and in discipline — the via negativa versus real-and-unreached — and so must not be collapsed.

Proof. Error (i) applies the projective discipline ("real, unreached, reachable by richer access") to the substrate, treating its lacking an axis-value as a fact awaiting measurement; but the substrate's lacking an axis-value is categorial, not a hidden fact (Articles 2, 4), so the error reifies. Error (ii) applies the apophatic discipline ("no fact, the question being ill-formed") to the floor, treating an unreached distinction as malformed; but the distinction is well-typed and determinate (Article 14), so the error is verificationist. Each error is the misapplication of one boundary's discipline at the other, and each is avoided by respecting the sources: the substrate is axis-free (no fact to find), the access is projective (real facts unreached). Hence the two must be kept distinct. □

The proposition locates the two great misreadings the volume is at pains to avoid, and shows them to be one mistake made in opposite directions. The reificationist hears that the substrate cannot be positively predicated and reads it as incompleteness — surely there is some deeper fact, some code or field, that better physics will reach — and so fills the apophatic silence with a projection installed as substrate. The verificationist hears that some distinctions can never be measured and reads it as meaninglessness — surely what cannot be measured is no fact — and so empties a real distinction into apophatic silence. Each takes one boundary for the other; each gets the discipline exactly backward. The volume's whole method, from the reification diagnosis of Article 2 to the incompleteness theorem of Article 14, is the steady refusal of both, and this article names the refusal as a single principle: respect the source of each silence, and never speak at one boundary in the voice of the other.

6.2 Counterarguments and Responses

Objection 1 (perhaps the substrate's axis-freeness is itself just our current incompleteness, to be overcome). *Response.* This is error (i) stated as a hope. Axis-freeness is not a measurement shortfall but a categorial fact: the substrate is what the axes are projections of, so it cannot itself bear an axis-value without circularity (a projection of itself), as Article 1 established. No future measurement changes a type; better physics resolves more distinctions *within* the projected world (lowering projective floors), but it cannot make an axis-predicate apply to the axis-free ground, any more than a sharper ruler measures the colour of a number. The hope misreads a type as a quantity.

Objection 2 (perhaps the projective floor's distinctions are, in the unreachable cases, genuinely factless, so verificationism is right there). *Response.* This is error (ii) stated as a concession. The floor's distinctions are, by the incompleteness theorem, genuine differences of the projected world — configurations $x \neq x'$ that a richer access would distinguish (Article 14); their reality does not depend on the present access's reaching them, only their measurement does. To make their reality depend on measurability is verificationism, and it is refuted by the lowerability of the floor: distinctions once unreachable have been reached by new projections throughout the history of physics, which would be unintelligible if the unreached were factless. Reality outruns a given access; that is why access can be enriched.

6.3 Section Result and Implications

Collapsing the two silences yields a symmetric double error — reification and verificationism — avoided by respecting each boundary's source and discipline (Proposition 6.1). The implication is that the volume's method is the steady refusal of both. We close by affirming the interior the boundaries bound.

7 The Interior Is Rich

7.1 Neither skeptical nor mystical

A volume that ends on two limits invites a misreading: that its lesson is the smallness of knowledge. The map forbids it.

Proposition 7.1 (The boundaries bound a rich interior)**.** *The interior bounded by the two silences — the recovered and projectable — is rich: it comprises every projection classified in Part II and every structure the model-builder recovers in Part III, the whole positive content of physics as the volume reads it. The two boundaries bound this interior and keep it honest; they do not empty it. The account is therefore neither skeptical (it affirms a vast knowable world) nor mystical (its one ineffability, the substrate's, is the explicable fact of axis-freeness). The limits are the edges of knowledge, not its measure.*

Proof. The interior's content is the union of Parts II and III: the six projections and their morphisms (Part II), and the model-builder's recovery, the arrow of time, and the relocated simulation (Part III), all well-typed and all within reach. That this content is vast and substantive is the demonstrated matter of those parts. The categorial boundary excludes only ill-typed predications of the substrate (Article 4); the projective boundary excludes only the joint fiber of a given access (Article 14); neither touches the interior, which they bound from outside and inside respectively. Skepticism would deny the interior's knowability, which the volume establishes; mysticism would make the substrate's ineffability a depth beyond understanding, which Section 4 reduced to axis-freeness. Hence neither, and a rich bounded interior. □

The proposition is the volume's guard against its own ending. Books that arrive at the limits of knowledge often leave the impression that the limits are the message — that the upshot is how little we can know, how the deepest things recede. The substrate frame ends otherwise. Its two limits are thin and well-understood: one is the edge where axis-predicates stop applying to the axis-free ground, the other the edge where a given access stops resolving real distinctions, and between them lies the entire knowable world, classified and recovered. The limits are not a verdict of ignorance but a discipline of honesty — they tell us precisely where not to speak positively (the substrate) and where not to fall silent (the floor) — and a discipline of honesty about the edges is what keeps the vast interior trustworthy. The volume's lesson is not that little can be known but that what can be known is bounded cleanly, and that knowing the bounds is part of knowing well.

7.2 Counterarguments and Responses

Objection 1 (an account whose ground is approached only by negation is mystical whatever it protests). *Response.* The via negativa here issues from a type distinction, not a mystery: the substrate is axis-free, axis-predicates apply only within axes, so they do not apply to it, and we say so by negation for that exact reason. A mystical apophasis holds the ground to exceed understanding; the present one holds it to be of the wrong type for axis-predicates, which is a matter of logic fully understood. The test is whether anything remains unexplained once the type is grasped, and nothing does: the negation is complete and its reason transparent. Structure shared with mysticism, mystery not.

Objection 2 (calling the interior "rich" is rhetoric; the limits are what is philosophically significant). *Response.* The limits are significant, and the article has given each a theorem and a discipline; the richness of the interior is not rhetoric but the demonstrated content of two-thirds of the volume, Parts II and III. Both matter, and the article's point is their relation: the limits are edges of the rich interior, not substitutes for it. To treat the limits as the whole significance is the very misreading Proposition 7.1 guards against — mistaking the edges for the country.

7.3 Section Result and Implications

The two boundaries bound a rich, knowable interior and keep it honest; the account is neither skeptical nor mystical (Proposition 7.1). The implication is that the volume's close is an affirmation bounded by discipline, not a retreat into limits. We record status and conclude.

8 Epistemic Status: Proved, Imported, Affirmed

In the manner of the series, we close with an explicit ledger. *Proved (here)* marks results established in this article; *Imported* marks results from prior articles; *Affirmed* marks positions taken; *Kinship (noted, distinguished)* marks the apophatic-tradition relation.

Claim	**Status**	**Locus**
Substrate axis-free; reification diagnosis	Imported	Kriger, IST 1 [2], Kriger, IST 2 [3]
Categorial (apophatic) boundary: ill-typed predicates	Imported	Kriger, IST 4 [4]
Projective boundary: incompleteness theorem, joint fiber	Imported	Kriger, IST 14 [7]
Maximal observer as non-reducing receiver	Imported	Kriger, IST 10 [5]
The reach has two boundaries; rich interior between	Proved (here)	Prop. 3.1
Via negativa: only negative/relational predication of substrate	Proved (here)	Prop. 4.1
Projective discipline: real, unreached, lowerable not closable	Proved (here)	Prop. 5.1
The double error (reification / verificationism)	Proved (here)	Prop. 6.1
The interior is rich; account neither skeptical nor mystical	Proved (here)	Prop. 7.1
Substrate's ineffability = axis-freeness (not mystery)	Affirmed	§4
Boundary of sense (locus classicus)	Affirmed (located)	Wittgenstein [8]
Limit-to-negation in the substrate frame	Imported (prior corpus)	Kriger [1]
Kinship to the apophatic method of older traditions	Kinship (noted, distinguished)	§4

The ledger marks the division of labor. The article *imports* the substrate's axis-freeness, the categorial boundary, the incompleteness theorem, and the maximal observer; it *proves* that the reach has two boundaries bounding a rich interior, that each demands its own discipline of speech, that collapsing them is a symmetric double error, and that the account is neither skeptical nor mystical; it *affirms* that the substrate's ineffability is axis-freeness; and it *notes and distinguishes* the kinship to the apophatic method, claiming the structure and declining the mystery.

9 Discussion

The result of this article is a unified account of the boundary of a subsystem's reach. The volume reached two limits by two routes: the categorial boundary at the axis-free substrate, where predicates fail to denote and questions are ill-formed, and the projective boundary within the projected world, where well-typed questions about real distinctions go unresolved by a given access. This article sets them side by side as the two edges of one reach, gives the map on which a rich interior — the recovered and projectable of Parts II and III — lies between them, and states the discipline of speech each edge demands. At the categorial edge the discipline is the via negativa: of the substrate one says what it is not and what it stands in relation to, never a positive axis-value, since every such filling reifies a projection into the substrate. At the projective edge the discipline is the opposite of silence: one asserts a determinate, real fact, unreached by the present access, lowerable by a richer access but never closable short of the non-projecting limit. The two errors to refuse are symmetric — reifying the substrate into a hidden fact, and emptying a real distinction into a non-fact — and each collapses the two silences into one. The account is neither skeptical nor mystical: the limits are thin, well-understood edges, and the interior they bound is the whole knowable world the volume has charted.

The article completes the bridge's turn toward the close. Part III opened with the central asymmetry — projection loses, the model recovers — and followed the recovering model through the arrow of time and the relocation of the simulation to the incompleteness of measurement, where it met the projective floor. This article joins that projective floor to the categorial boundary of the foundations, so that the two limits the volume found at its two ends — the apophatic silence of Article 4 and the incompleteness theorem of Article 14 — are seen as the two edges of one reach, each with its discipline, bounding the interior the rest of the volume built. What remains is the volume's final synthesis, drawing the whole into one view. The substrate frame began by refusing to mistake a projection for the substrate; it ends by mapping exactly where speech must turn negative and where it must not fall silent, and by affirming, between those two edges, a world that can be known.

9.1 Limitations

Three limitations are material. *First*, the article presupposes the categorial boundary of Article 4 and the incompleteness theorem of Article 14; a reader who rejects either will not accept the unified account, and the article re-argues neither. *Second*, the disciplines of speech are stated as principles — the via negativa, the threefold projective stance — and illustrated, but their application to specific contested cases (which predications of a given physical posit are axis-internal, which distinctions lie in a given floor) is not worked through and would require case-by-case analysis. *Third*, the kinship to the apophatic method of older traditions is noted and distinguished but not studied; a serious comparison, which the author's other work pursues elsewhere, lies outside the scope of a physics article and is not attempted here.

9.2 Future Directions

The immediate sequel is the volume's closing synthesis, which can now draw the foundations, the projections, the bridge, and the two boundaries into a single statement of the substrate frame. Beyond the volume, the unified account invites a precise question continuous with the series: whether the two boundaries are related by the inter-axis morphisms of Article 3 — whether the categorial boundary is the limit of the morphism diagram as the axes are stripped away and the projective floor its limit as access is held fixed — so that both edges would be read off one structure. That question, joining the apophatic and the incomplete in a single geometry of the knowable, is left to the work that can address it; the present article's task was to map the two silences, and the map is drawn.

Acknowledgments

The author thanks colleagues of the Information Physics Institute and the Institute of Integrative and Interdisciplinary Research for sustained discussion of the foundations of the series.

References

[1] B. Kriger. The law of the limit to negation: on the impossibility of absolute nothing and the necessity of a differentiated ground. IIIR. Zenodo, 2026. doi:10.5281/zenodo.18101306.

[2] B. Kriger. The information substrate: differentiation as the ground of information. IST, Volume VII, Article 1. Zenodo, 2026.

[3] B. Kriger. Substrate and projection: why the universe needs no information. IST, Volume VII, Article 2. Zenodo, 2026.

[4] B. Kriger. The categorial boundary: on predicates that do not cross the block. IST, Volume VII, Article 4. Zenodo, 2026.

[5] B. Kriger. The axis of the receiver: Bateson's difference and the pragmatic projection. IST, Volume VII, Article 10. Zenodo, 2026.

[6] B. Kriger. Projection loses, the model recovers: the central asymmetry and the opening of the bridge. IST, Volume VII, Article 11. Zenodo, 2026.

[7] B. Kriger. The incompleteness of measurement: the residual floor as a theorem of projective access. IST, Volume VII, Article 14. Zenodo, 2026.

[8] L. Wittgenstein. *Tractatus Logico-Philosophicus.* Kegan Paul, Trench, Trubner & Co., London, 1922.

A Code for Figure 1

Figure 1 is a schematic map drawn with nested rounded rectangles: the innermost the reached interior, the surrounding ring the well-typed unreached floor, the outer frame the projected world (its border the categorial boundary), and the background the axis-free substrate. The two boundaries — projective (inner) and apophatic (outer) — are labeled with their disciplines of speech. No numerical computation is involved; the figure encodes the structure of Propositions 3.1, 4.1, and 5.1.

The Substrate Frame Entire: One Descent from Differentiatedness to the Boundaries of Reach

The volume as a single argument, turning on one distinction and one discipline, and the puzzles it dissolves at a stroke

Boris Kriger[1,2]

[1] Information Physics Institute
boris.kriger@informationphysicsinstitute.net

[2] Institute of Integrative and Interdisciplinary Research, Toronto
boriskriger@interdisciplinary-institute.org

ORCID: https://orcid.org/0009-0001-0034-2903

The Information Substrate Theory (IST) — Volume VII, Article 16
Part III: The Bridge — the synthesis

Abstract

This article is the keystone of the volume: it draws the foundations, the projections, and the bridge into a single statement of the substrate frame and shows that the whole is one argument turning on one distinction. The descent runs without a gap. Something exists, since absolute nothing is impossible; what exists is necessarily differentiated; information is the name of that differentiatedness, not a substance or a fluid; the axes along which the world is read — time, space, energy — and the laws stated in their terms are facets, projections of the differentiated substrate, which is itself axis-free and carries no projection value. The known measures of information are then projections along six axes, classified and related by their morphisms and shown to materialize across the scales of the prior volumes. Physical projection loses irreversibly, and a model-building subsystem recovers the projectable structure by inference rather than by inversion; from that central asymmetry follow the arrow of time as the subsystem's record gradient, the relocation of the simulation into the model, the incompleteness of measurement as a theorem of projective access, and the two boundaries that bound a subsystem's reach. We show that this entire descent turns on a single distinction — substrate versus projection — and a single discipline: never mistake a projection for the substrate. From the one distinction the whole follows, and under it a family of standing puzzles dissolves together rather than one at a time: whether the universe computes, whether the world is made of information, whether it is a simulation, what the arrow of time is, why measurement cannot be complete. These are not five problems with five solutions but one confusion — projection taken for substrate — corrected once. Finally we say plainly what the frame is and is not: an ontology and a discipline, conservative about physics, which it affirms entirely; not a new physics, adding no equation; not a speculative metaphysics, positing no hidden entity; and neither idealism nor mysticism. The volume is one argument with one premise and one discipline, bounding one knowable world.

Keywords: information substrate; substrate and projection; differentiation; synthesis; reification; dissolution of puzzles; ontology; foundations of physics.

Contents

1 Introduction

A volume of seventeen articles risks reading as seventeen arguments, and this one, the keystone, exists to show that it is one. The substrate frame was built in three movements — foundations, projections, bridge — and each movement was developed in its own articles with its own results; but the movements are not independent contributions gathered under a title. They are the stages of a single descent, each entailed by the one before, from a single premise to a bounded conclusion. This article states the descent whole, exhibits the one distinction on which it all turns, shows the puzzles it dissolves to be one confusion corrected once, and says plainly what the frame is and is not.

The descent runs without a gap, and it is worth tracing in a breath. Something exists, because absolute nothing is impossible [3]. What exists is necessarily differentiated, for an undifferentiated existent would be no existent at all. Information is the name of that differentiatedness — not a substance the world is made of, not a fluid it contains, but the word for its being differentiated [3]. The axes along which we read the world — time, space, energy — and the laws we state in their terms are facets of the differentiated substrate, projections of it onto single coordinates; the substrate itself is axis-free and carries no projection value, no number of bits, no entropy, no program length [3, 4]. The known measures of information are accordingly projections: Shannon's along the channel axis, Kolmogorov's along the description axis, and the rest, descending from the substrate by the projection theorem, classified into six axes and related by their morphisms, and materializing across the scales of the prior volumes from the nucleon to the cosmic web [5, 7, 8]. Physical projection loses irreversibly — the data-processing inequality, the second law, and the uncomputability of inversion are its three faces — and a model-building subsystem recovers the projectable structure not by inverting the projection, which is blocked, but by inference [9]. From that central asymmetry the bridge unfolds: the arrow of time is the subsystem's gradient of accumulated records [10]; the simulation, real as the limit of the recovering model, lives in the subsystem and not in the world [11]; the incompleteness of measurement is the non-trivial joint fiber of any proper projective access [12]; and two boundaries, the apophatic at the substrate and the projective within the world, bound a subsystem's reach around a rich interior [13].

The whole of this turns on one distinction. Substrate is not projection; the differentiated ground is not any of the single-coordinate readings of it; and the one discipline the frame demands is never to mistake the second for the first. Every result of the volume is a consequence of holding that line. The known measures are projections because they are readings along axes, not the axis-free ground (Part II). The reification puzzles dissolve because each takes a projection — a computation, an information content, a simulation — for the substrate (Section 5). The model-builder's situation follows because a subsystem with projective access loses and recovers and is bounded exactly as projection dictates (Part III). One distinction, held steadily, generates the classification, dissolves the puzzles, and bounds the knowable; that is the unity this article exhibits.

And because the distinction is one, the puzzles it dissolves are one. It can seem that the frame answers a list of separate questions — does the universe compute, is the world information, is it a simulation, what is time's arrow, why is measurement incomplete — but they are not a list. They are a single confusion, projection taken for substrate, appearing in five costumes, and the frame does not solve them one by one; it removes the confusion once, and they go together. Seeing this is seeing the frame's economy: a single corrective move, not a battery of separate solutions. We end by stating, in the plainest terms, what the frame therefore is — an ontology and a discipline, conservative about physics, which it affirms entirely — and what it is not: not a new physics, not a speculative metaphysics, neither idealism nor mysticism. The volume is one argument, with one premise and one discipline, bounding one knowable world.

2 Logical Structure of the Argument

The article states the descent, isolates its hinge, shows the puzzles to be one, and characterizes the frame.

Step 1 (Section 3): The descent is one argument.
The stages from differentiatedness to the two boundaries form a single chain, each entailed by the prior (Proposition 3.1, Figure 1).

Step 2 (Section 4): One distinction, one discipline.
The whole turns on substrate versus projection and the discipline of not conflating them (Proposition 4.1).

Step 3 (Section 5): The puzzles are one confusion.
The standing puzzles are one reification appearing in several costumes, dissolved together (Proposition 5.1).

Step 4 (Section 6): What the frame is, and is not.
An ontology and a discipline, conservative about physics; not a new physics, not speculative metaphysics, not idealism or mysticism (Proposition 6.1).

Step 5 (Section 7): Epistemic status.
A ledger consolidates the volume's results by part and status.

The dependencies are those of a synthesis: Step 1 assembles the descent the prior articles built; Step 2 names its hinge; Step 3 draws the consequence for the puzzles; Step 4 places the whole. The thesis is that the substrate frame is one argument turning on one distinction, dissolving one confusion, and amounting to an ontology and a discipline rather than a new physics or a metaphysics.

3 The Descent Is One Argument

3.1 The single chain

We state the descent as a chain and claim its links established.

Proposition 3.1 (The volume is one argument)**.** *The substrate frame is a single chain of entailments, each link established in the cited articles:*

(1) *Absolute nothing is impossible, so something exists, and what exists is differentiated; information is the name of differentiatedness [3].*

(2) *Axes (time, space, energy) and the laws stated in them are facets; the substrate is axis-free and carries no projection value [3, 4].*

(3) *The known information measures descend from the substrate as projections along axes (projection theorem); some predicates of the substrate are ill-typed (apophatic boundary) [5, 6].*

(4) *There are six projection axes, related by morphisms and materialized across scales [7, 8].*

(5) *Physical projection loses irreversibly; the model-building subsystem recovers by inference, not inversion [9].*

(6) *Hence the arrow of time (record gradient), the relocation of the simulation, the incompleteness of measurement, and the two boundaries [10–13].*

Each link is entailed by the prior together with the cited result; the volume is therefore one argument, not a collection.

Remark in lieu of proof. The links are not re-proved here — each is the established content of its articles — but their *order of entailment* is the claim, and it is linear. Link (2) follows from (1): if information is the name of differentiatedness, the measures of information are readings of it along chosen coordinates, i.e. projections, and the ground read is axis-free. Link (3) follows from (2): the known measures, being projections, descend from the substrate, and predicates internal to an axis do not apply to the axis-free ground. Link (4) refines (3) into the classification. Link (5) follows because a projection is many-to-one, hence lossy and non-invertible, so recovery must be inferential. Link (6) follows because a subsystem with lossy projective access has a record asymmetry (arrow), a recovering model whose limit is simulation, a joint fiber (incompleteness), and two boundaries. Each step uses only its predecessors and one cited result; the chain is unbroken. □

The proposition and Figure 1 make visible what a volume's worth of separate articles can obscure: that there is a single thread. One does not adopt the foundations and

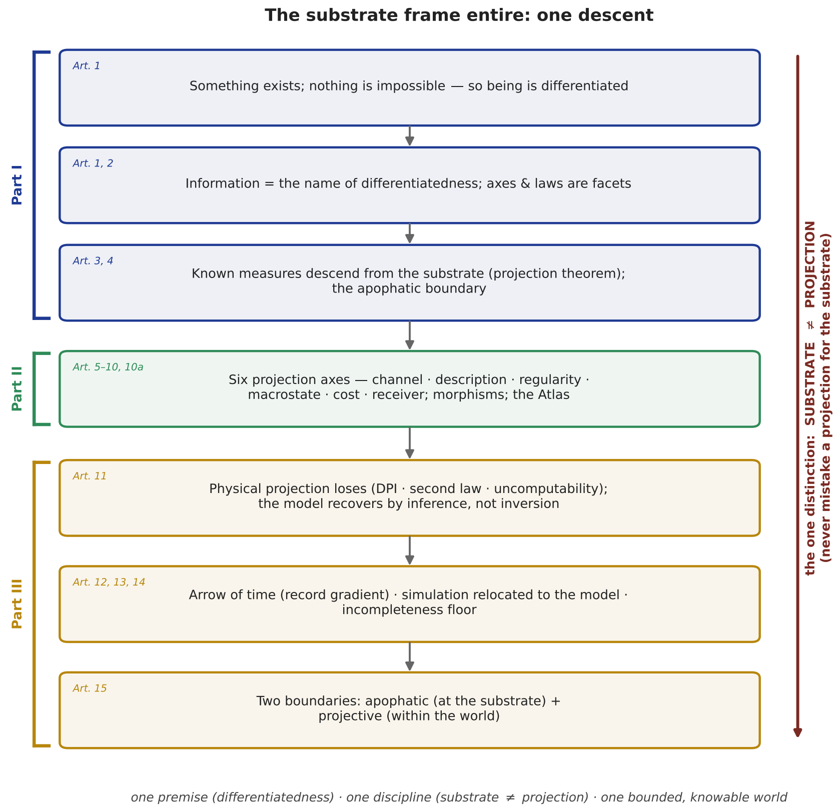

Figure 1: The substrate frame as one descent. The spine runs from the foundations (Part I, blue) through the projections (Part II, green) to the bridge and the two boundaries (Part III, amber), each stage entailed by the one above and annotated with its articles. Down the right runs the through-line that ties every stage together: the one distinction, substrate $\neq$ projection, and the discipline of never mistaking a projection for the substrate. One premise — differentiatedness — and one discipline yield one bounded, knowable world.

then, as a further and optional step, the projections, and then, as a third, the bridge; the projections are what the foundations make of the known measures, and the bridge is what the projections make of a subsystem that has only projections to go on. The architecture is not modular but sequential, a descent in which each level is the previous level's consequence for the next question. That is why the volume can be stated in a breath, as the introduction stated it: the breath follows the single thread.

3.2 Counterarguments and Responses

Objection 1 (the links are looser than "entailment"; each article adds substantive assumptions). *Response.* Each article adds apparatus — definitions, a theorem, a figure — but the *load-bearing* move at each link is the entailment stated, and the apparatus serves it. The projection theorem (link 3) does real work, but it works because the measures are projections (link 2); the incompleteness theorem (link 6) does real work, but it works because access is projective (link 5). The claim is not that no article adds anything, but that the spine is a chain of entailments, with each article's apparatus discharging its link. The objection is right that articles add content and wrong that the additions break the chain.

Objection 2 (a reader may accept some links and reject others, so the unity is overstated). *Response.* A reader may, and the chain shows exactly what such a reader must give up. Reject link (1) and nothing starts; accept (1) but reject (2)'s axis-freeness and the projections lose their ground; accept the projections but reject (5)'s lossiness and the bridge does not follow. The unity is not that the chain cannot be broken but that it is a chain — breaking a link drops everything below it. That is a stronger claim than a loose family of theses, and it is the one the proposition makes: the volume stands or falls in segments, from the top, not article by independent article.

3.3 Section Result and Implications

The volume is a single chain of entailments from differentiatedness to the two boundaries, each link established in its articles (Proposition 3.1, Figure 1). The implication is that the chain has a hinge. We name it.

4 One Distinction, One Discipline

4.1 The hinge of the whole

The descent is long, but it turns on a single point, repeated at every link.

Proposition 4.1 (One distinction, one discipline)**.** *The entire frame turns on the distinction between the* substrate *— the differentiated, axis-free ground — and a* projection *— a reading of it along a single axis. The one discipline the frame demands is to hold this distinction: never to predicate of the substrate what belongs to a projection, never to install a projection as the substrate. Every result is a consequence of holding the line: the measures are projections (not the ground); the substrate carries no projection value; physical projection loses because it is a many-to-one reading; the model recovers a projection's preimage, not the substrate unmediated; the boundaries mark where projection meets the axis-free (apophatic) and where projective access meets its fiber (projective). Drop the distinction and the frame collapses into the reifications it was built to avoid.*

Proof. Each result invokes the distinction. "The measures are projections" (Part II) is the distinction applied to information theory: a measure is a reading along an axis, not the ground. "The substrate carries no projection value" (Article 2) is the distinction stated of the ground. "Projection loses" (Article 11) is the many-to-one character of a reading, not of the ground, which loses nothing because it is not a reading. "The model recovers a preimage" (Article 11) recovers projectable structure, a projection's fiber, not the substrate. The two boundaries (Article 15) are the two places the distinction bites: the apophatic, where a projection's predicate is wrongly sought of the ground; the projective, where a subsystem's readings leave a fiber. Remove the distinction — identify substrate with projection — and each result becomes a reification: the ground *is* a measure, *is* a computation, *is* information. The discipline of the distinction is thus necessary and sufficient for the frame. □

The proposition identifies the frame's economy. A theory that needed a separate principle for each result would be a list; the substrate frame needs one, applied repeatedly. The single distinction — ground versus reading, substrate versus projection — is doing all the work, at every link of the descent, and the single discipline of not conflating them is the whole method. This is why the frame can be taught in a sentence and why its errors are all one error: to forget the distinction anywhere is to reify somewhere, and to hold it everywhere is to have the frame entire. The volume's length is the unfolding of one short idea, not the accumulation of many.

4.2 Counterarguments and Responses

Objection 1 (a single distinction cannot carry a whole physics-adjacent framework; something is being smuggled). *Response.* Nothing is smuggled; the distinction's power is exactly that it is a *type* distinction, and type distinctions are load-bearing precisely because they apply everywhere at once. Distinguishing a function from its argument, or a set from its members, organizes vast tracts of mathematics from one stroke; the substrate–projection distinction does the same for the reading of the world. It carries the framework not by hidden content but by being the right cut, repeated. That a single correct distinction organizes much is the ordinary power of getting a type right, not a smuggling.

Objection 2 (if the distinction is so simple, why seventeen articles). *Response.* Because applying a simple distinction correctly across information theory, thermodynamics, cosmology, cognition, and the philosophy of measurement requires showing, in each domain, what the distinction implies and what reifications it forbids — work that is neither short nor automatic. The distinction is simple; its disciplined application is not, and the articles are that application. A simple principle with intricate consequences is the normal shape of a foundational idea; the simplicity of the principle and the length of its working-out are not in tension.

4.3 Section Result and Implications

The frame turns on one distinction — substrate versus projection — and one discipline, holding it; every result follows, and dropping it collapses the frame into reification

(Proposition 4.1). The implication is that the puzzles the frame dissolves are one confusion. We show it.

5 The Puzzles Are One Confusion, Dissolved Once

5.1 Five costumes, one error

The frame is often met as a set of answers to a set of questions. The questions are one.

Proposition 5.1 (The standing puzzles are one reification)**.** *Several standing puzzles about information and the world are a single confusion — a projection taken for the substrate — in different costumes, and the frame dissolves them together:*

(i) Does the universe compute / process information? *It need not: computation is a projection we apply, not an activity of the ground (Article 2).*

(ii) Is the world made of information ("it from bit")? *No: information is the name of differentiatedness, not a substance the world is made of (Article 1).*

(iii) Is the world a simulation? *The question, on the substrate reading, is ill-typed: it installs a computational projection as a substrate beneath the world (Article 13).*

(iv) What is the arrow of time? *The model-builder's record gradient, not a flow of the substrate, which is a-temporal (Article 12).*

(v) Why can measurement never be complete? *Because projective access has fibers; the limit is the joint fiber, not a failing of instruments (Article 14).*

Each is the confusion of a projection (a computation, an information-substance, a simulation, a flow, a complete reading) with the substrate or with what projective access can be. The frame does not solve five problems; it removes one confusion, and the five dissolve together.

Proof. Each item is the distinction of Proposition 4.1 applied to a costume of the one error. *(i)* treats a description-axis projection (computation) as the ground's activity; corrected, the universe need not compute (Article 2). *(ii)* treats information as a substance rather than the name of differentiatedness; corrected, "made of information" is a category error (Article 1). *(iii)* treats a computation as a substrate beneath the world; corrected, the reading is ill-typed (Articles 4, 13). *(iv)* treats the time axis's gradient as a flow of the a-temporal ground; corrected, the arrow is the subsystem's (Articles 1, 12). *(v)* treats the limit of measurement as contingent rather than as the joint fiber of projective access; corrected, incompleteness is structural (Article 14). In each, the same correction — do not take a projection for the substrate, or for more than projective access can be — dissolves the puzzle; so they are one confusion corrected once. □

The proposition is the frame's most visible payoff and its clearest economy. The puzzles it addresses are among the most discussed at the border of physics and philosophy, and they are usually treated as distinct, each with its own literature and its own proposed answers. The substrate frame declines to add a sixth answer to each; it observes that the five are one question wearing five costumes — in every case, a projection has been promoted to the substrate or mistaken for the whole of what can be accessed — and that correcting the one promotion answers all five. This is not a claim to have solved five hard problems by cleverness but a claim that they were never five: that a single confusion, once named and refused, takes its several appearances with it. The economy is the test of the distinction's being the right one.

5.2 Counterarguments and Responses

Objection 1 (dissolving puzzles rather than solving them is evasion; the questions are real). *Response.* The questions are real as questions people ask, and the frame takes them seriously enough to diagnose precisely why each misfires — which costume of the reification it wears — rather than dismissing them. Dissolution is not evasion when it is earned: it shows the question to rest on a confusion and shows the confusion exactly. Where a well-formed residue remains — a subsystem can build a simulation, measurement has a determinate floor — the frame affirms it (Articles 13, 14). What it declines is to answer the malformed part as though it were well-formed, which would be the evasion of pretending a confusion is a discovery.

Objection 2 (perhaps some of the five are genuinely distinct and only resemble one another superficially). *Response.* They differ in subject matter — computation, substance, simulation, time, measurement — and resemble one another not superficially but structurally: each, when analyzed, is found to promote a projection to the substrate or to mistake projective access for unmediated access (the proof traces the promotion in each). The claim is not that the five are the same question on the surface but that they share one underlying error, which the analysis exhibits case by case. Different surfaces, one structure, is precisely what makes the single correction dissolve them all.

5.3 Section Result and Implications

The standing puzzles — computation, it-from-bit, simulation, the arrow, incompleteness — are one reification in five costumes, dissolved together by the one distinction (Proposition 5.1). The implication is that the frame's contribution is a corrective ontology, not a new physics. We say what it is.

6 What the Frame Is, and Is Not

6.1 An ontology and a discipline

A synthesis should say plainly what has been built. We state it, with its limits.

Proposition 6.1 (What the substrate frame is and is not)**.** *The substrate frame is an* ontology *— a claim about what there is at bottom: a differentiated, axis-free substrate, of which the axes and their measures are projections — and a* discipline *— a rule for speaking of it: never mistake a projection for the substrate. It is* conservative about physics*: it adds no equation, predicts no new number, and affirms physical theory entirely, re-describing the status of its quantities rather than revising them. It is* not *a new physics;* not *a speculative metaphysics positing hidden entities, the substrate being the differentiatedness physics already studies, not an extra thing behind it;* not *idealism, since the projected world and its substrate are affirmed as real and mind-independent; and* not *mysticism, since the substrate's only ineffability is its axis-freeness, fully explicable.*

Proof. That the frame is an ontology and a discipline is the content of Parts I and II: an account of the ground and a rule for predication [3, 4]. That it is conservative about physics is shown by Article 2, which affirms that the universe needs no information added to its physics, and by Article 10a, which re-describes established results (the nucleon, the cosmic web) as materializations without revising them [4, 8]; the frame changes the *status* of physical quantities (projections, not substances) and leaves their *content* untouched. It is not metaphysics in the speculative sense because it posits no entity beyond what physics studies — the substrate is differentiatedness, not a hidden thing (Article 1). It is not idealism because the recovering model is answerable to a mind-independent world (Article 13). It is not mysticism because the apophatic silence reduces to axis-freeness (Article 15). Each negative is the explicit content of the cited article. □

The proposition fixes the frame's modest and exact ambition, against two opposite misreadings. The first overstates it: the frame is taken for a new fundamental physics, a rival to or replacement of established theory, from which novel predictions should flow. It is not; it adds no equation and affirms the physics entire, and its contribution is to the *interpretation* of physical quantities — their status as projections — not to their content. The second understates it: the frame is taken for mere words, a relabeling that changes nothing. It is more than that, for getting the type of the world's measures right dissolves a family of puzzles and disciplines a family of inferences, which words that changed nothing could not do. Between overstatement and understatement lies the frame's actual claim: an ontology and a discipline, conservative about physics, that re-describes the known and dissolves the confused, and does so by one distinction held without exception.

6.2 Counterarguments and Responses

Objection 1 (an interpretation that yields no new prediction is idle; only predictions matter). *Response.* Interpretation that dissolves standing puzzles and disciplines inference is not idle, even without new predictions: it changes which questions are pursued and which are seen to misfire, which is consequential for inquiry. Much foundational work — the interpretation of probability, of spacetime, of measurement — yields clarity rather than new numbers and is not therefore idle. The frame's value is of this kind, and it is honest about it: it claims clarity and discipline, not prediction, and a reader who counts only predictions is applying a standard the frame does not pretend to meet, not exposing a failure to meet its own.

Objection 2 (calling the substrate "differentiatedness physics already studies" makes it redundant — either it is physics or it is a posit). *Response.* It is neither redundant nor a posit beyond physics: it is the *differentiatedness* of the world physics studies, considered as such, prior to its projection onto any axis. Physics studies the projections — fields, states, entropies — and the substrate is what those are projections of, named and characterized (Articles 1, 2). That is not an extra entity (which would be a speculative posit) nor a mere synonym for physics (which would be redundant); it is the ground of the projections physics works with, made explicit. The frame's ontology is the explicitation of what physics presupposes, not an addition to it.

6.3 Section Result and Implications

The substrate frame is an ontology and a discipline, conservative about physics; not a new physics, not speculative metaphysics, not idealism, not mysticism (Proposition 6.1). The implication is that the volume's contribution is clarity disciplined by one distinction. We record status and close.

7 Epistemic Status: the Volume Consolidated

In the manner of the series, we close with a ledger, here consolidating the volume by part and status rather than by single claim. *Established (in the volume)* marks results proved across the articles; *Imported (prior corpus)* marks results from the broader programme the volume rests on; *Conjectured* marks the open horizon; *Stance* marks the disciplines the frame adopts.

Claim	Status	Locus
Differentiatedness as ground; information as its name	Established (in the volume)	Kriger, IST 1 [3]
Substrate axis-free; carries no projection value	Established (in the volume)	Kriger, IST 1 [3], Kriger, IST 2 [4]

continued on next page

Consolidated ledger (continued)

Claim	**Status**	**Locus**
Known measures descend as projections (projection theorem)	Established (in the volume)	Kriger, IST 3 [5]
Six axes; morphisms; Atlas materialization	Established (in the volume)	Kriger, IST 5 [7], Kriger, IST 10a [8]
Projection loses; model recovers by inference	Established (in the volume)	Kriger, IST 11 [9]
Arrow, simulation relocated, incompleteness, two boundaries	Established (in the volume)	Kriger, IST 12 [10], Kriger, IST 13 [11], Kriger, IST 14 [12], Kriger, IST 15 [13]
The volume is one chain of entailments	Established (here)	Prop. 3.1
One distinction (substrate $\neq$ projection) carries all	Established (here)	Prop. 4.1
The standing puzzles are one reification, dissolved once	Established (here)	Prop. 5.1
Frame is ontology + discipline, conservative about physics	Stance	Prop. 6.1
Impossibility of absolute nothing; differentiation as condition	Imported (prior corpus)	Kriger [1, 2]
Morphism geometry; H_∞ as projective invariant	Conjectured	§8

The consolidated ledger is the volume in one view: the established descent across Parts I–III, the three synthetic claims of this article (one chain, one distinction, one dissolved confusion), the frame's stance, the prior corpus it rests on, and the conjectural horizon it leaves open. What the volume establishes it marks established; what it imports it marks imported; what it leaves to future work it marks conjectured; and what it adopts as discipline it marks stance, not theorem.

8 Discussion

The result of this article is the volume seen whole. The substrate frame is one descent: from the impossibility of absolute nothing to the differentiated ground, from differentiatedness to information as its name, from there to the axes as facets and the known measures as projections, from the lossiness of projection to the recovering model, and from the model-builder's situation to the two boundaries that bound its reach. The descent is a single chain, each link the prior link's consequence, and it turns at every link on one distinction — substrate versus projection — held by one discipline: never mistake a projection for the substrate. From that distinction the classification of the projections follows, the model-builder's arrow and simulation and incompleteness follow, and the two silences follow; and under it the standing puzzles of the field — whether the universe computes, whether the world is information, whether it is a simulation, what time's arrow is, why measurement is incomplete — are seen to be one reification in several costumes, dissolved not one at a time but together. The frame is, in the end, an ontology and a discipline: an account of the differentiated ground physics presupposes, and a rule for speaking of it without reification. It is conservative about physics, which it affirms entirely, and it is neither a new physics nor a speculative metaphysics nor idealism nor mysticism.

What remains is the volume's final article, which can now stand back from the frame and ask what it leaves open and where it sits in the wider programme. For the synthesis has not closed every question; it has, if anything, sharpened the open ones. The deepest is the conjecture the series has raised at each stage: that the inter-axis morphisms of the projection theorem form a geometry from which both the unity of the measures and the limits of their recovery could be read — the completeness of the classification and the incompleteness floor as two faces of one structure. If that geometry exists, the substrate frame would acquire a mathematical spine to match its conceptual one, and the descent this article traced in prose would have a diagram. Whether it exists is not settled here. What is settled is the frame itself: one premise, one distinction, one discipline, one bounded and knowable world, set out in a single descent that this article has, at its keystone, drawn together.

8.1 Limitations

Three limitations are material. *First*, the synthesis presupposes the results it consolidates; it re-proves nothing, and a reader who rejects a link of the descent rejects everything below it, as Proposition 3.1 makes explicit. *Second*, the claim that the standing puzzles are one confusion is argued by exhibiting the shared structure case by case, not by a general theorem that all such puzzles reduce to reification; puzzles outside the five treated would need their own analysis. *Third*, the characterization of the frame as conservative about physics holds for the volume's re-description of established results; whether every future physical posit can be accommodated as a projection without strain is not guaranteed and would be tested case by case.

8.2 Future Directions

The immediate sequel is the volume's final article, standing back to place the frame and mark its open horizon. Beyond the volume, the central open question is the morphism geometry: whether the inter-axis maps of Article 3 form a structure whose invariants are the unity of the measures and whose limits are the incompleteness floor of Article 14, so that classification and incompleteness would be two faces of one geometry. Establishing or refuting that geometry would be the natural continuation of the programme, giving the substrate frame the mathematical articulation its conceptual descent invites. That work is left to what can address it; the present article's task was to draw the frame together, and the frame is drawn together.

Acknowledgments

The author thanks colleagues of the Information Physics Institute and the Institute of Integrative and Interdisciplinary Research for sustained discussion of the foundations of the series, and acknowledges the prior volumes of the programme on which this synthesis rests.

References

[1] B. Kriger. Testing nothingness: on the impossibility of absolute non-being and the necessity of a differentiated ground. IIIR. Zenodo, 2026. doi:10.5281/zenodo.18764754.

[2] B. Kriger. Differentiation as the ontological condition of actualization. IIIR. Zenodo, 2026. doi:10.5281/zenodo.18268520.

[3] B. Kriger. The information substrate: differentiation as the ground of information. IST, Volume VII, Article 1. Zenodo, 2026.

[4] B. Kriger. Substrate and projection: why the universe needs no information. IST, Volume VII, Article 2. Zenodo, 2026.

[5] B. Kriger. The projection theorem: how known information measures descend from the substrate. IST, Volume VII, Article 3. Zenodo, 2026.

[6] B. Kriger. The categorial boundary: on predicates that do not cross the block. IST, Volume VII, Article 4. Zenodo, 2026.

[7] B. Kriger. The axis of the channel: Shannon information as a substrate projection. IST, Volume VII, Article 5. Zenodo, 2026.

[8] B. Kriger. The Atlas projection: where the substrate axes materialize across the scales of the SPCA programme. IST, Volume VII, Article 10a. Zenodo, 2026.

[9] B. Kriger. Projection loses, the model recovers: the central asymmetry and the opening of the bridge. IST, Volume VII, Article 11. Zenodo, 2026.

[10] B. Kriger. The arrow of time: how the model-builder experiences the block. IST, Volume VII, Article 12. Zenodo, 2026.

[11] B. Kriger. The relocation of the simulation: why simulation lives in the model, not in the world. IST, Volume VII, Article 13. Zenodo, 2026.

[12] B. Kriger. The incompleteness of measurement: the residual floor as a theorem of projective access. IST, Volume VII, Article 14. Zenodo, 2026.

[13] B. Kriger. The two silences: the apophatic and projective boundaries of a subsystem's reach. IST, Volume VII, Article 15. Zenodo, 2026.

A Code for Figure 1

Figure 1 is a schematic drawn with stacked rounded rectangles forming a vertical spine: seven stages from the foundations through the projections to the bridge and the two boundaries, grouped by part (brackets at left) and tied by the through-line at right (substrate $\neq$ projection). No numerical computation is involved; the figure encodes the chain of Proposition 3.1 and the hinge of Proposition 4.1.

The Lens and the Open Geometry: The Place of the Substrate Frame and the Problem It Leaves Open

Volume VII as the lens that reads the Atlas, the morphism geometry as the central conjecture, and the frame as a portable discipline

Boris Kriger[1,2]

[1] Information Physics Institute
boris.kriger@informationphysicsinstitute.net

[2] Institute of Integrative and Interdisciplinary Research, Toronto
boriskriger@interdisciplinary-institute.org

ORCID: https://orcid.org/0009-0001-0034-2903

The Information Substrate Theory (IST) — Volume VII, Article 17
Part III: The Bridge — the close of the volume

Abstract

The synthesis stated the substrate frame whole; this closing article stands back from it to place it, mark what it leaves open, and say how it is meant to be used. Three things are done. First, we place Volume VII within the programme of which it is the seventh part. The prior six volumes built an Atlas of structure across the scales of the cosmos — the nucleon, the atom, the knotted molecule, the living cell, the cosmic web — by physical methods; Volume VII is the lens that reads that Atlas, the information-theoretic ontology under which each scale is the materialization of a projection. The six are the realizations and the seventh is the frame, and together they form one programme, monolithic not because every scale is the same but because each realizes a facet of one differentiated substrate. Second, we state the central open problem the volume bequeaths. The projection theorem related the axes by inter-axis morphisms, and the incompleteness theorem found a residual floor; we conjecture that a single geometry — the diagram of the axes and their morphisms — carries both, with the completeness of the classification as the diagram's connectedness and the incompleteness floor as the substrate region no morphism into an access covers. If the conjecture holds, classification and incompleteness are two faces of one structure, and the conceptual descent of the volume acquires a mathematical spine. We state the conjecture precisely and mark it, throughout, as a conjecture. Third, we offer the frame as a portable discipline rather than a system to be believed: a single diagnostic question — does this predicate of the substrate what belongs only to a projection? — that travels beyond the volume and disciplines reasoning about information wherever the temptation to reify arises. We close with the restraint the series has kept: the frame claims an ontology, a discipline, and a set of dissolutions, and it refuses to be a theory of everything, a source of predictions, or a creed. The volume ends not with a flourish but with a clear statement of what was built, what remains, and how to carry it.

Keywords: substrate frame; research programme; morphism geometry; completeness; incompleteness; portable discipline; reification; foundations of physics.

Contents

1 Introduction

The synthesis drew the substrate frame together into a single descent and a single distinction [5]. With the frame stated whole, the volume's last article has a different task: not to build, but to place, to mark, and to hand over. To place the frame within the programme of which it is one part; to mark the central problem it leaves unsolved; and to hand it over as something usable — a discipline a reader can carry, not only a system a reader can assent to. These three are the work of the close.

The placement first. Volume VII did not arrive from nowhere. It is the seventh part of a programme whose prior six volumes built, by the ordinary methods of physics and a topological–dissipative analysis across scales, an Atlas of structure in the cosmos: the nucleon as a self-consistent closure, the atom as a contraction to its spectrum, the knotted topology of biomolecules, the deviation of the living cell from its attractor, the cosmic web mapped by survey data, the closed topology of the universe [3]. Those volumes computed and measured; Volume VII interprets. It is the lens that reads the Atlas — the ontology under which each computed object is seen as the materialization, at its scale, of a projection of one differentiated substrate. The relation was exhibited in the Atlas article [3]; here we state it as the architecture of the whole: six volumes of realizations and one of the frame, forming a single programme. The programme is monolithic, but not because it claims every scale to be the same thing; it is monolithic because each scale, in its own physics, realizes a facet of the one substrate, and the seventh volume is the account of that one substrate and its facets.

The central open problem next. The volume proved two things that ask to be joined. The projection theorem related the six axes by inter-axis morphisms — rescaling, kernel-averaging, refinement — and conjectured the classification complete [2]. The incompleteness theorem found, for any proper projective access, a residual floor: the joint fiber no measurement resolves [4]. We conjecture that these are two faces of one geometry. The diagram whose nodes are the axes and whose edges are the morphisms would carry, in its connectedness, the completeness of the classification — every known measure an object, every relation a composite — and would locate, in its complement, the incompleteness floor — the substrate region no morphism into a given access reaches. If the conjecture holds, the volume's conceptual descent acquires a mathematical spine, and what the classification preserves and what the floor discards are read off one structure. We state the conjecture carefully and, as the series has done throughout, mark it a conjecture and not a result.

The handing-over last. The frame is offered not as a creed to be believed but as a discipline to be practiced, and a discipline travels. Its portable core is a single diagnostic question, asked of any claim about information and the world: does this predicate of the substrate something that belongs only to a projection? Where the answer is yes, a reification is in progress, and the frame says how to correct it. The question outlasts the volume; it applies to claims the volume never treated, wherever the temptation to install a projection as the ground arises. We end with the restraint the series has kept: the frame claims an ontology, a discipline, and a set of dissolutions, and refuses to be a theory of everything, a generator of predictions, or a system demanding assent. The volume closes with a plain statement of what was built, what remains open, and how to carry what was built into questions still to come.

2 Logical Structure of the Argument

The article places the frame, states the open problem, offers the discipline, and closes.

Step 1 (Section 3): The place of Volume VII.
The frame is the lens reading the Atlas the prior volumes built; the seven volumes form one programme (Proposition 3.1).

Step 2 (Section 4): The central open problem.
A single morphism geometry would carry both the completeness of the classification and the incompleteness floor (Conjecture 4.1, Figure 1).

Step 3 (Section 5): The frame as a portable discipline.
A single diagnostic question carries the frame beyond the volume (Proposition 5.1).

Step 4 (Section 6): Claim and restraint.
The frame claims an ontology, a discipline, and dissolutions, and refuses to be a theory of everything, a source of predictions, or a creed (Proposition 6.1).

Step 5 (Section 7): Epistemic status.
A ledger marks the placement, the conjecture, and the stance.

The dependencies are those of a close: Step 1 places what was built, Step 2 marks what remains, Step 3 hands it over, Step 4 states the spirit in which it is offered. The thesis is that Volume VII is the lens of a seven-volume programme, that its central open problem is the morphism geometry, and that the frame is best carried as a discipline rather than believed as a system.

3 The Place of Volume VII: the Lens and the Atlas

3.1 Six realizations and one frame

We state the architecture of the programme of which this volume is the seventh part.

Proposition 3.1 (Volume VII is the lens that reads the Atlas)**.** *The prior volumes of the programme built an Atlas of structure across the scales of the cosmos by physical methods — the nucleon and atom at the microscale, the knotted molecule and living cell at the mesoscale, the cosmic web and the closed cosmic topology at the macroscale. Volume VII is the lens that reads this Atlas: the information-theoretic ontology under* *which each Atlas object is the materialization, at its scale, of a projection of one differentiated substrate (Article 10a). The six prior volumes are the realizations and the seventh is the frame; together they are one programme, monolithic not because every scale is the same but because each realizes a facet of the one substrate.*

Remark in lieu of proof. That the Atlas was built by the prior volumes and that each object materializes a projection is the content of the Atlas article, which mapped the description axis to the nucleon and atom, the regularity axis to the knotted molecule and cell, the channel and macrostate axes to the cosmic web, and the global register to the cosmic topology [3]. The present proposition states the relation as architecture: the prior volumes supply the realizations, by physics; this volume supplies the frame, the ontology of substrate and projection under which the realizations are read as facets of one ground. The monolithic character is the Atlas article's anti-reification result restated — the unity is that each scale realizes a facet, not that the substrate is any measure — and so the programme coheres without collapsing its scales into one. □

The proposition fixes how the volume should be read in its setting. A reader who came to Volume VII alone might take it for a free-standing philosophy of information; a reader who knows the programme sees it as the keystone of an arch whose other stones are the computed and measured structures of the prior volumes. The frame is not speculation laid over physics but the lens through which the programme's own physics is seen as one — the account that lets the nucleon and the cosmic web, computed by entirely different means at scales forty orders of magnitude apart, be recognized as materializations of the same few projections of the same differentiated ground. The seventh volume earns its place not by adding to the Atlas but by reading it, and the programme is complete as an arch only with the keystone in.

3.2 Counterarguments and Responses

Objection 1 (a lens that adds no physics is dispensable; the Atlas stands without it). *Response.* The Atlas stands as physics without the frame, and the frame insists on exactly that (it is conservative about physics). What the frame adds is not physics but unity and discipline: the recognition that the Atlas's scales realize a few common projections, and the refusal to reify those projections into the substrate. That recognition is dispensable for computing any single scale and indispensable for seeing the scales as one programme; a reader content with the scales severally may omit it, a reader who wants the programme whole may not. The lens is optional for the parts and constitutive of the whole.

Objection 2 (calling the programme "monolithic" overstates a loose family of results across disparate fields). *Response.* The monolithic claim is precise and modest: each scale realizes a facet of one substrate, established case by case in the Atlas article, not that the fields are reducible to one another or that one theory predicts all scales. The unity is at the level of the frame — one ground, several projections, materialized severally — not at the level of the physics, which remains plural and scale-specific. "Monolithic" names the unity of the frame over plural realizations, which is exactly what the Atlas article established and the synthesis consolidated, neither more nor less.

3.3 Section Result and Implications

Volume VII is the lens that reads the Atlas built by the prior volumes; the seven form one programme, six realizations and one frame (Proposition 3.1). The implication is that the programme, like the volume, leaves a central problem open. We state it.

4 The Central Open Problem: the Morphism Geometry

4.1 The conjecture

The volume proved two results that ask to be joined, and joining them is the central open problem it bequeaths.

Conjecture 4.1 (The morphism geometry). *Let D be the diagram whose objects are the six projection axes and whose morphisms are the inter-axis maps of the projection theorem — rescaling ($S = k_B H$), kernel-averaging ($\mathbb{E}[K] \approx H$), refinement ($K = K_{\text{eff}} + R$), and the relations of the receiver and cost axes (the maximal observer, the Maxwell demon). Then:*

(A) (Completeness.) *Every known information measure is an object of D, and every relation between measures is a composite of its morphisms: the classification is complete, and D is connected [2].*

(B) (Incompleteness.) *For any proper projective access $\mathcal{P}$, the residual floor $H_\infty(\mathcal{P})$ is computable from D as the entropy of the substrate region not in the image of any morphism into $\mathcal{P}$ — the diagram's complement relative to the access [4].*

If the conjecture holds, the completeness of the classification (Part II) and the incompleteness of measurement (Article 14) are two faces of one geometry: what D connects, and what lies outside its image.

The conjecture is stated as a conjecture, and the emphasis is deliberate. The volume has been careful, at every stage, to mark what it proves and what it surmises, and the morphism geometry is the largest of its surmises — the structure that would unify the classification and the incompleteness floor into one mathematical object, giving the conceptual descent of the synthesis a formal spine. We do not establish it. We state it

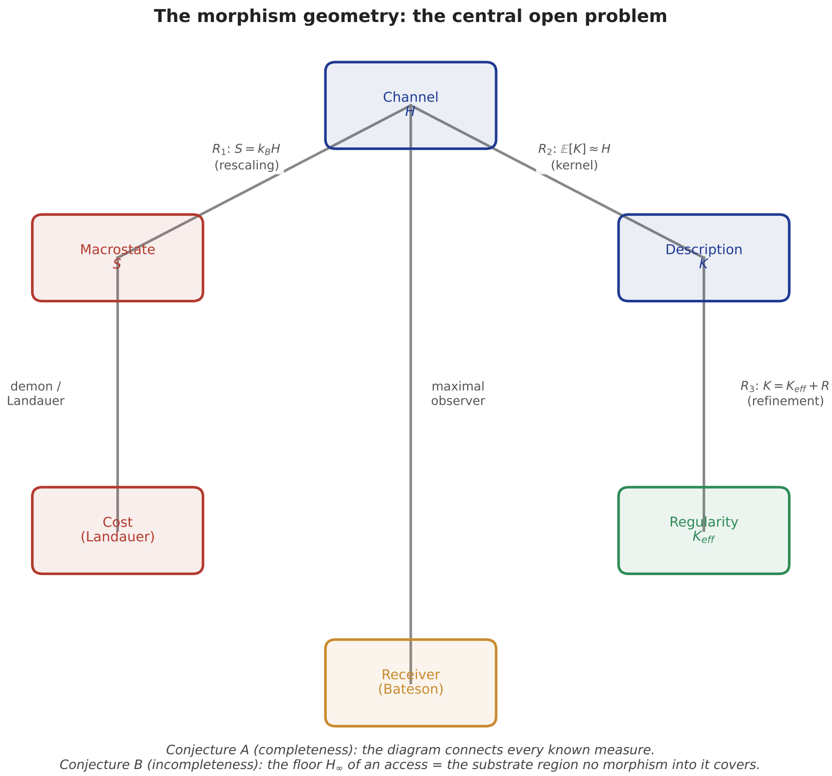

Figure 1: The morphism geometry as the central open problem. The six projection axes are nodes; the inter-axis morphisms of the projection theorem are edges — rescaling (R_1), kernel-averaging (R_2), refinement (R_3), the maximal-observer relation of the receiver axis, and the demon/Landauer relation of the cost axis. The conjecture is that this diagram carries both faces of the volume: its connectedness the completeness of the classification (Conjecture 4.1A), and its complement relative to an access the incompleteness floor (Conjecture 4.1B). The figure states the problem; it does not claim it solved.

precisely enough to be a research target: a diagram of axes and morphisms (Figure 1), a completeness claim about its connectedness, and an incompleteness claim about its complement. Whether the diagram is genuinely a category with the conjectured properties, whether the floor is genuinely its complement, and whether the morphisms suffice to connect every measure — these are open, and they are the natural continuation of the programme. The conjecture's appeal is its economy: if true, the two great results of the volume, one from Part II and one from Part III, would be the image and the complement of a single structure.

4.2 Counterarguments and Responses

Objection 1 (the conjecture is too vague to be a research target; "the diagram's complement" is not defined). *Response.* The conjecture is stated at the level of precision appropriate to a closing article — objects, morphisms, a connectedness claim, a complement claim — and it is explicitly a conjecture, not a theorem awaiting only details. Making "the diagram's complement relative to an access" fully precise — as, say, the cokernel of the morphisms into the access, or the substrate quotient by their joint image — is itself part of the open problem, and different precisifications are worth trying. The article's task is to pose the problem clearly enough to be worked on, not to solve it; a posed problem with alternative precisifications is a research target, not a vagueness.

Objection 2 (treating information measures as objects of a category may simply not work). *Response.* It may not, and the conjecture is falsifiable in that direction: if some known measure is not an object of D, or some relation is not a composite of the morphisms, completeness fails; if the floor is not recoverable from the diagram, the incompleteness face fails. That the conjecture can fail is a virtue, marking it a genuine claim rather than a slogan. The projection theorem already established three of the morphisms with their composition relations [2], so the category-theoretic reading is not idle speculation but the extrapolation of a started construction; whether it completes is exactly the open question.

4.3 Section Result and Implications

The central open problem is the morphism geometry: whether one diagram carries both the completeness of the classification and the incompleteness floor, as image and complement (Conjecture 4.1, Figure 1). The implication is that the volume bequeaths a precise target, not merely a mood. We turn from what remains to how the frame is carried.

5 The Frame as a Portable Discipline

5.1 One diagnostic question

The frame is most useful not as a system to assent to but as a discipline to practice, and a discipline reduces to a habit.

Proposition 5.1 (The frame travels as a diagnostic question)**.** *The substrate frame's portable core is a single diagnostic question, applicable to any claim about information and the world:* does this predicate of the substrate something that belongs only to a projection? *Where the answer is no, the claim is in order. Where it is yes, a reification is in progress, and the frame's correction applies: relocate the predicate to the projection* *where it belongs, and leave the substrate axis-free. The question outlasts the volume and applies to claims the volume never treated; it is the frame in usable form, a habit rather than a creed.*

Argument. The whole frame turns on the substrate–projection distinction and the discipline of not conflating them (Article 16). To practice that discipline on a given claim is exactly to ask whether the claim predicates of the substrate what belongs to a projection — the diagnostic question. If it does not, no conflation occurs; if it does, the conflation is the reification Article 2 diagnoses, corrected by relocating the predicate to its projection. Since the distinction is general, the question is general: it applies wherever a claim about information and the world might conflate ground and reading, including claims outside the volume's cases. The frame thus reduces, for use, to one repeatable question. □

The proposition is the article's gift to the reader who will not adopt the whole frame but might use part of it. One need not accept the descent from the impossibility of nothing, nor the classification of six axes, nor the bridge, to find the diagnostic question useful; one need only suspect, of some confident claim that the universe *is* information, or *is* a computation, or *is* a simulation, that a reading has been mistaken for the ground, and ask the question. The frame's value travels at the level of the question even when its system stays behind. This is the sense in which the volume offers a discipline rather than a doctrine: a doctrine is believed whole or not at all, a discipline is practiced wherever it helps, and the substrate frame is built to be practiced — one question, asked whenever the temptation to reify a projection arises.

5.2 Counterarguments and Responses

Objection 1 (a diagnostic question detached from the system is just a slogan against information talk). *Response.* It is not a blanket prohibition on information talk; it is a test that passes most such talk and flags only the reifying kind. "The signal carries two bits," "the genome encodes control information," "the web shows neighbour mutual information" all pass — they predicate of projections, correctly. Only claims that predicate of the *substrate* what belongs to a projection — the world *is* information, *is* a computation beneath physics — are flagged. The question discriminates, which a slogan does not; it is the system's discipline in portable form, with the system available behind it for anyone who wants the justification.

Objection 2 (without the full system, a user cannot tell substrate-predication from projection-predication). *Response.* The full system is exactly what supplies the criterion, and it is available; the portable question is a shortcut for the practiced, not a replacement of the justification for the unpracticed. A user uncertain whether a predicate is axis-internal can consult the classification (Part II) and the categorial boundary (Article 4), which say which predicates belong to which axes and which fail to apply to the substrate. The question is the handle; the volume is the tool behind it. Offering a handle does not deny the tool.

5.3 Section Result and Implications

The frame travels as one diagnostic question — does this predicate of the substrate what belongs only to a projection? — usable beyond the volume, a discipline rather than a creed (Proposition 5.1). The implication is that the frame is offered to be practiced, in a spirit the close should state. We state it.

6 Envoi: Claim and Restraint

6.1 What the volume claims, and refuses

A volume should end by saying, without inflation, what it has and has not claimed.

Proposition 6.1 (The frame's claim and its restraint)**.** *The substrate frame claims three things and refuses three. It* claims *an ontology (a differentiated, axis-free substrate of which the known measures are projections), a discipline (never mistake a projection for the substrate), and a set of dissolutions (the reification puzzles, resolved together). It* refuses *to be a theory of everything (it adds no physics and predicts no new number), a generator of predictions (its contribution is interpretive and disciplinary), and a creed (it asks practice of a discipline, not assent to a system). The volume ends, accordingly, with a statement of what was built and what remains, not with a claim to have finished.*

Remark in lieu of proof. The three claims are the established content of the volume: the ontology of Part I [1], the discipline that runs throughout [5], and the dissolutions of the synthesis [5]. The three refusals are the volume's explicit self-characterization: conservative about physics, hence no theory of everything and no new predictions (Article 16), and offered as a discipline, hence no creed (Section 5). The restraint is not modesty for its own sake but accuracy: to claim more — a final theory, a predictive engine, a doctrine — would be to reify the frame's reach, the very error the frame forbids of the substrate. The close is therefore disciplined by the frame's own rule. □

The proposition closes the volume in the register it has kept. The series has been, throughout, an exercise in saying exactly as much as is warranted and no more — marking proofs as proofs, conjectures as conjectures, analogies as analogies, and reifications as errors. It would be untrue to that register to end with a flourish, a claim that the nature of information has been settled or the puzzles of physics dissolved for good. What has been done is bounded and statable: an ontology of the differentiated substrate and its projections has been set out, a discipline for speaking of them without reification has been established and applied, a family of standing puzzles has been shown to be one confusion and dissolved, and a central problem — the morphism geometry — has been posed and left open. That is the volume, and to say more would be to do the thing the volume warns against. The frame ends by declining to overreach, which is the last and most fitting application of its own discipline.

6.2 Counterarguments and Responses

Objection 1 (a framework that predicts nothing and claims to be no theory of everything is too modest to matter). *Response.* Modesty about scope is not smallness of consequence. The frame dissolves puzzles that have occupied the border of physics and philosophy for decades, disciplines inference about information across fields, and unifies a seven-volume programme under one lens — consequences a predictive theory does not supply and a clarifying frame does. The refusal of prediction is honesty about the kind of contribution, not a confession of triviality; foundational clarity and predictive power are different goods, and the frame claims the first plainly rather than pretending to the second.

Objection 2 (ending on an open conjecture leaves the volume unfinished). *Response.* The volume is finished as what it is — an ontology, a discipline, a set of dissolutions — and open as what it points toward, the morphism geometry. A foundational work that closed every question it raised would be suspect, not admirable; the honest end of such a work is a clear account of what is settled and a precise statement of what is not, which is what the close provides. The conjecture is not a loose end but a bequeathed problem, marked as such. The volume ends complete in its claims and open in its horizon, which is the proper shape of a foundational close.

6.3 Section Result and Implications

The frame claims an ontology, a discipline, and a set of dissolutions, and refuses to be a theory of everything, a source of predictions, or a creed (Proposition 6.1). The implication is that the volume ends as it proceeded: saying exactly what is warranted. We record the final status.

7 Epistemic Status: the Close

In the manner of the series, we close with a final ledger. *Established (in the volume)* marks the volume's proven content; *Placement* marks the architectural claim of this article; *Conjectured* marks the open problem; *Stance* marks the spirit of the offering.

Claim	Status	Locus
The substrate frame (ontology, discipline, dissolutions)	Established (in the volume)	Kriger, IST 16 [5]
Atlas materialization of the axes across scales	Established (in the volume)	Kriger, IST 10a [3]
Inter-axis morphisms of the projection theorem	Established (in the volume)	Kriger, IST 3 [2]
Incompleteness floor of projective access	Established (in the volume)	Kriger, IST 14 [4]
Volume VII is the lens reading the Atlas; one programme	Placement	Prop. 3.1
The morphism geometry (completeness + incompleteness as one)	Conjectured	Conj. 4.1
The frame travels as one diagnostic question	Established (here)	Prop. 5.1
The frame's claim (ontology, discipline, dissolutions)	Stance	Prop. 6.1
The frame's refusal (not a theory of everything / predictions / creed)	Stance	Prop. 6.1

The final ledger is the close in one view: the volume's established content, the placement of the volume within the programme, the conjecture it bequeaths, the discipline it hands over, and the stance in which it is offered. What the volume proved it marks established; what this article asserts about the volume's place it marks placement; what it leaves to future work it marks conjectured; and the spirit of the whole it marks stance, not theorem.

8 Discussion

The result of this article is the close of the volume. Volume VII is the lens of a seven-volume programme: the prior six built an Atlas of structure across the scales of the cosmos by physical means, and the seventh reads that Atlas as the materialization, scale by scale, of the projections of one differentiated substrate. The programme is one, monolithic not because its scales collapse into one thing but because each realizes a facet of the one ground. The volume bequeaths a central open problem, stated as a conjecture: that a single morphism geometry carries both the completeness of the classification and the incompleteness of measurement, the former in its connectedness and the latter in its complement, so that what the volume proved in Part II and what it proved in Article 14 would be two faces of one structure. And the volume hands itself over as a discipline rather than a creed, reducible for use to a single diagnostic question — does this predicate of the substrate what belongs only to a projection? — that travels wherever the temptation to reify a projection arises. The frame claims an ontology, a discipline, and a set of dissolutions, and refuses to be a theory of everything, a source of predictions, or a doctrine demanding assent.

So the volume ends where its discipline directs. It set out to do for the foundations of information what the prior volumes did for the cosmos: not to add a new substance or a new law, but to get the type of things right — to see that information is the name of differentiatedness, that the known measures are projections, that the substrate is axis-free, and that almost every puzzle at the border of physics and philosophy about information is a single confusion of the reading for the read. Having shown this, having classified the projections, traced the bridge from loss to recovery, and mapped the two silences that bound a subsystem's reach, the volume has done its work. What remains — the morphism geometry, the further anchoring across scales, the application of the diagnostic question to questions not yet posed — is the programme's continuing horizon, and it is left, with a precise statement of what it requires, to the work that can address it. The substrate frame is set out, placed, and handed over; the volume is closed.

8.1 Limitations

Three limitations are material. *First*, the placement of Volume VII within the programme presupposes the prior volumes and the Atlas article; a reader without them has the frame but not its realizations, and the monolithic claim rests on work outside this volume. *Second*, the morphism geometry is a conjecture, stated at the precision of a closing article and not established; its central terms admit alternative precisifications, and whether any yields the conjectured unity is open. *Third*, the diagnostic question is offered as a portable discipline, but its correct application in contested cases requires the full system behind it, so the portability is real for the practiced and partial for the newcomer.

8.2 Future Directions

The programme's continuation is the morphism geometry: to make the diagram of axes and morphisms a precise mathematical object, to test whether it is complete in the conjectured sense, and to determine whether the incompleteness floor is its complement relative to an access. Beyond that, the diagnostic discipline invites application to questions the volume did not treat — the status of information in quantum foundations, in biology, in the theory of computation — wherever a projection risks being installed as a substrate. And the programme as a whole invites the further anchoring of the frame across the scales of the Atlas, extending the materialization of the axes as new structure is computed and measured. These are the horizon; the volume's task was to set out the frame, place it, and hand it over, and that task is complete.

Acknowledgments

The author thanks colleagues of the Information Physics Institute and the Institute of Integrative and Interdisciplinary Research for sustained discussion of the foundations of the series, and acknowledges the six prior volumes of the programme that this volume reads, and the readers who have followed the argument to its close.

References

[1] B. Kriger. The information substrate: differentiation as the ground of information. IST, Volume VII, Article 1. Zenodo, 2026.

[2] B. Kriger. The projection theorem: how known information measures descend from the substrate. IST, Volume VII, Article 3. Zenodo, 2026.

[3] B. Kriger. The Atlas projection: where the substrate axes materialize across the scales of the SPCA programme. IST, Volume VII, Article 10a. Zenodo, 2026.

[4] B. Kriger. The incompleteness of measurement: the residual floor as a theorem of projective access. IST, Volume VII, Article 14. Zenodo, 2026.

[5] B. Kriger. The substrate frame entire: one descent from differentiatedness to the boundaries of reach. IST, Volume VII, Article 16. Zenodo, 2026.

[6] B. Kriger. *Volume I: The Local Gravitation of Quantum Vacuum — A Unified Solution to the Dark Sector (αLGQV Theory Monograph)*. IIIR Cosmology and Theoretical Physics, Toronto, 2026. doi:10.5281/zenodo.19027460.

[7] B. Kriger. *Volume VI: Structural Persistence and Coherence Analysis — A Topological–Dissipative Method Across Physical Scales*. IIIR Cosmology and Theoretical Physics, Toronto, 2026. doi:10.5281/zenodo.20455316.

A Code for Figure 1

Figure 1 is a schematic graph: six nodes for the projection axes and five edges for the inter-axis morphisms of the projection theorem (rescaling, kernel-averaging, refinement, the maximal-observer relation, and the demon/Landauer relation). No numerical computation is involved; the figure encodes the diagram D of Conjecture 4.1 and poses, as its caption states, the open problem of whether D carries both the completeness of the classification and the incompleteness floor.