

# The Operational Structure of Physical Distinguishability: Reframing, Convergence, and Hierarchy at the Causal-Patch Bound

VERSF Theoretical Physics Program

---

## For the General Reader

The number  $10^{123}$  has appeared in physics before. It is the holographic entropy of the observable universe — the maximum information content of the region of space we can observe, calculated from the surface area of the cosmic horizon in units set by quantum gravity. Physicists have known this number exists for decades.

What has not been done — and what this paper does — is ask three sharper questions about it.

The first question is: what does this number actually bound? The existing literature says it bounds *entropy* — a measure of disorder or information content in the thermodynamic sense. This paper argues that it bounds something more fundamental and more physically meaningful: the number of *operationally distinguishable states*, meaning states that any physical process could, even in principle, tell apart. These two things sound similar but are not the same, and the difference matters enormously. Moving from entropy to distinguishability connects the bound directly to measurement theory and to what it means for a physical difference to be real, rather than merely mathematical.

The second question is: three independent lines of argument — a thermodynamic optimisation, four established physical constraints, and a pre-geometric derivation from the void-fold framework — all arrive at the same bound. Is that a coincidence? This paper shows it is not, and makes that non-coincidence precise. When three frameworks with no shared premises all produce the same structural constraint, that convergence is itself evidence that the constraint reflects something deep about the architecture of physical reality.

The third question is: the holographic bound, the Lloyd computational bound, and the Landauer thermodynamic bound all bear on the same number but give slightly different values —  $10^{123}$ ,  $10^{120}$ , and  $10^{90}$  respectively. Are these three competing estimates, or three layers of one coherent structure? This paper proposes that they are the latter: a hierarchy in which each layer is the relevant bound for a different class of physical question. Making that hierarchy precise is both a clarification of what has already been established and a new analytical tool for asking which bound applies in which context.

Together, these three contributions represent the original theoretical work surrounding the causal-patch distinguishability limit. They do not depend on any claim that the number itself is

new. They depend on asking, and answering, questions about its structure that have not previously been asked.

---

## Abstract

The causal-patch holographic distinguishability limit  $\mathcal{L} \approx 2.3 \times 10^{123}$  — the quantity the companion paper *Taylor's Number: The Computational Boundary of Physical Reality* derives through thermodynamic optimisation and labels  $\mathcal{L}_T$  — is numerically identical to the holographic entropy bound of the observable universe's causal horizon. This is not an approximation: the companion paper explicitly sets  $\mathcal{L}_T = \pi(R_U/\ell_P)^2$  in its appendices, equating Taylor's Number with the holographic bound as its final defined value. The number therefore predates the VERSF programme. The originality of the programme's treatment of this quantity does not rest on the number but on three structural contributions that the existing holographic literature has not developed.

**Contribution I — Operational Reframing.** The holographic literature bounds entropy. The present paper formally develops the distinct conceptual move of treating  $\mathcal{L}$  as a bound on *operationally distinguishable states* — states separable by a physical process requiring finite energy, time, and entropy production. This reframing connects the holographic bound to measurement theory and to a precise criterion for what constitutes a physical difference, a connection the existing literature does not make explicit.

**Contribution II — Three-Route Convergence.** A thermodynamic optimisation (companion paper), four established empirical constraints (companion paper II), and a pre-geometric void-fold derivation (synthesis paper) independently arrive at the same structural conclusion. We formalise this convergence, demonstrate the mutual independence of the three routes, and argue that the convergence is not forced by dimensional analysis alone. The convergence is itself a structural result about the architecture of physical reality.

**Contribution III — Distinguishability Hierarchy.** The holographic, Lloyd, and Landauer bounds give numerically different values —  $10^{123}$ ,  $10^{120}$ , and  $\sim 10^{90}$  respectively — which the existing literature treats as loosely consistent estimates. We show these are not competing approximations but three distinct layers of a single nested hierarchy: the holographic bound as geometric ceiling on state-space capacity, the Lloyd bound as dynamical instantiation limit on configurations physically realised over cosmic history, and the Landauer bound as operational access floor on distinctions thermodynamically resolvable at current conditions. Each layer is the appropriate bound for a different class of physical claim.

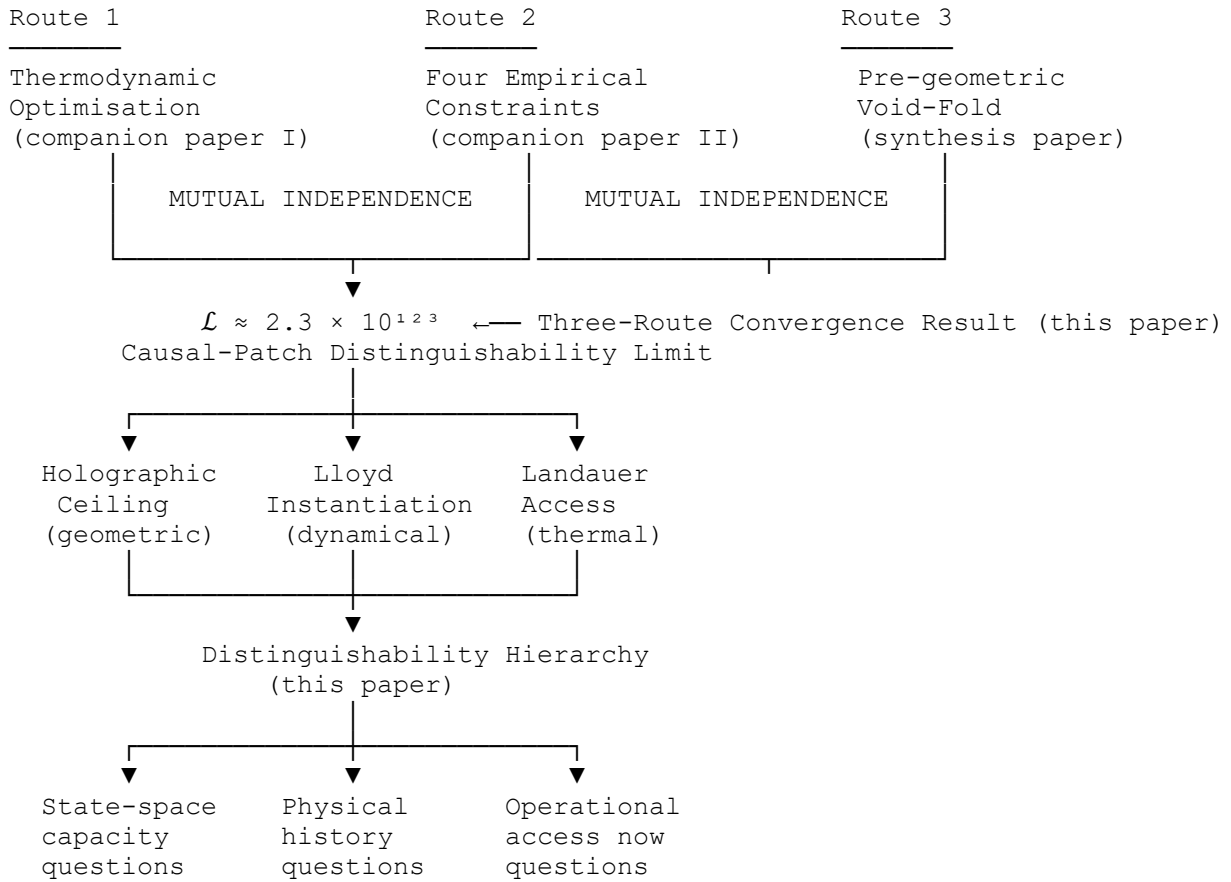
**Keywords:** physical distinguishability, operational definition, holographic bound, convergence argument, distinguishability hierarchy, Lloyd bound, Landauer principle, causal patch, VERSF programme

---

# Table of Contents

- Programme Context
  - 1. Introduction: Three Questions About a Known Number
    - 1.1 What Is and Is Not Claimed to Be New
    - 1.2 On the Decision to Name the Quantity
    - 1.3 Why These Questions Have Not Been Asked
    - 1.4 Structure of the Paper
  - 2. Contribution I: The Operational Reframing
    - 2.1 What the Holographic Literature Actually Bounds
    - 2.2 The Conceptual Gap: From Entropy to Distinguishability
    - 2.3 The Operational Definition as Formal Bridge
    - 2.4 What the Reframing Changes
    - 2.5 Connection to Measurement Theory
  - 3. Contribution II: The Three-Route Convergence
    - 3.1 The Three Routes and Their Independence
    - 3.2 Formalising the Convergence Claim
    - 3.3 Why the Convergence Is Non-Trivial
    - 3.4 The Convergence as Structural Evidence
  - 4. Contribution III: The Distinguishability Hierarchy
    - 4.1 The Three Bounds and Their Values
    - 4.2 The Hierarchy Proposal
    - 4.3 Layer 1: The Holographic Ceiling
    - 4.4 Layer 2: The Lloyd Instantiation Limit
    - 4.5 Layer 3: The Landauer Access Floor
    - 4.6 The Nested Structure and Its Implications
  - 5. Consequences: What Each Contribution Implies
    - 5.1 Consequences of the Operational Reframing
    - 5.2 Consequences of the Convergence Result
    - 5.3 Consequences of the Hierarchy
  - 6. Open Questions
  - 7. Conclusion
  - References
-

# Programme Context



## 1. Introduction: Three Questions About a Known Number

### 1.1 What Is and Is Not Claimed to Be New

The quantity  $\mathcal{L} \approx 2.3 \times 10^{123}$  is the holographic entropy bound of the observable universe's causal patch. In the form  $S_{\max} = A / 4\ell_P^2$ , applied to the cosmic horizon area, it appears in the work of Bekenstein (1973), was given its holographic interpretation by 't Hooft (1993) and Susskind (1995), and was extended to dynamical spacetimes by Bousso (1999, 2002). As a number, it predates the VERSF programme entirely.

It is important to be precise about the relationship between Taylor's Number and the holographic entropy bound, because this relationship determines where the originality of the VERSF programme actually lies. The companion paper *Taylor's Number: The Computational Boundary of Physical Reality* defines  $\mathcal{L}_T$  via the raw dimensionless scale ratio  $(R_U/\ell_P)^2 \approx 7.2 \times 10^{121}$  and then, in its appendices, sets  $\mathcal{L}_T = I_{\max} = \pi(R_U/\ell_P)^2 \approx 2.3 \times 10^{123}$  — explicitly equating Taylor's Number with the holographic entropy bound as its final value. The two quantities are therefore numerically identical by the companion paper's own definition. Taylor's Number is not

a different number that happens to be close to the holographic bound. It is the holographic bound, given a new name and a new conceptual framing.

This makes the originality question precise. It cannot rest on any numerical distinction between Taylor's Number and the holographic entropy bound, because there is none. It rests entirely on what the companion paper does with that number: it reframes the holographic entropy bound as a bound on physically distinguishable states rather than on thermodynamic entropy, derives it through a thermodynamic optimisation argument rather than from black hole thermodynamics, and draws consequences for computation, measurement, and mathematical meaningfulness that the holographic literature has not developed. Whether those moves constitute original contributions is the question this paper addresses.

This paper makes no claim that  $\mathcal{L}$  is a new number. The claim, stated precisely, is that three structural questions about  $\mathcal{L}$  have not been developed together as a unified framework in the existing literature, and that addressing them in combination constitutes original theoretical work.

The three questions are:

**Q1.** The holographic literature bounds entropy. Does bounding entropy bound distinguishability? Are these the same thing, and if not, what is the precise bridge between them?

**Q2.** Three independent theoretical frameworks — a thermodynamic optimisation, a set of four established empirical constraints, and a pre-geometric derivation — all arrive at  $\mathcal{L}$  as a structural limit. Is this convergence coincidental, and if not, what does it imply?

**Q3.** The holographic, Lloyd, and Landauer bounds give numerically distinct values —  $10^{123}$ ,  $10^{120}$ , and  $\sim 10^{90}$  — for ostensibly related quantities. Are these competing approximations to the same thing, or distinct layers of a hierarchy? If the latter, which layer applies to which class of physical question?

These three questions are logically independent of each other and each has a determinate answer. Providing those answers is the paper's contribution.

## 1.2 On the Decision to Name the Quantity

A natural question arises immediately: if  $\mathcal{L}$  is numerically the holographic entropy bound, already present in the literature under that name, why give it a new name at all? Is "Taylor's Number" not simply a renaming — a piece of branding that adds nothing to the physics?

The question deserves a direct answer, because it goes to the heart of what naming does in theoretical physics and whether it is ever legitimate to name a known quantity.

**Names track conceptual roles, not just numerical values.** Physical constants and limits are not named merely because they exist. They are named when a specific value becomes the natural unit of a specific conceptual domain — when a number plays a structurally distinguished role in a framework such that referring to it repeatedly without a name becomes cumbersome and the

unnamed state obscures the argument. When a known number takes on a new and specific conceptual role in a new framework, naming it for that role is a way of signalling to the reader that the number is doing particular work in this framework, not merely appearing as a background fact.

The parallel for the present case is this. The holographic entropy bound of the causal patch has existed in the literature as a bound on entropy in the context of black hole thermodynamics and quantum gravity. What the Taylor's Number companion paper does is give that number a new conceptual role: as the boundary of operationally distinguishable physical states, as the natural unit of the finite distinguishability framework, and as the quantity around which the optimisation, empirical, and pre-geometric arguments converge.

**The naming is prospective, not retrospective.** In standard physics practice, constants are named after individuals only retrospectively, by the community, after the work has been accepted and the contribution recognised. Naming one's own constant in the proposing paper is a departure from this norm, and it carries risks: it can appear presumptuous, and it invites the objection — pursued in this paper — that the quantity is not new. We acknowledge these risks directly.

The justification offered here is not that the naming is costless or unambiguous, but that it is defensible on the following grounds. The reframing of the holographic entropy bound as a distinguishability limit is a substantive conceptual contribution, not merely a notational preference. The three structural contributions of the present paper — the operational bridge, the convergence argument, and the hierarchy — are contributions to what this number *means* and *does*, not to its numerical value. When a conceptual reframing is substantive enough to change what questions a quantity answers and what theoretical work it performs, naming the quantity for that reframing is a way of marking the reframing as a distinct intellectual act. It is also honest: it makes clear that the VERSF programme is making a specific claim about the significance of this number in a new framework, rather than obscuring the relationship to prior work.

**What the name should not be taken to imply.** The name "Taylor's Number" should not be taken to imply that the numerical value is new, that the holographic literature was unaware of the quantity, or that the derivation via thermodynamic optimisation is independent of all prior work. All three of these implications would be false, and the present paper is designed precisely to make that clear. What the name marks is the specific conceptual role the quantity plays within the VERSF programme: as the operational bound on physical distinguishability, as the convergence point of three independent theoretical routes, and as the quantity whose hierarchical structure — geometric ceiling, dynamical instantiation limit, operational access floor — the programme develops for the first time.

If the community ultimately concludes that the reframing is not sufficiently original to warrant a distinct name, that judgement is appropriate and the physics stands independently of the naming. The bound, the bridge, the convergence, and the hierarchy are contributions to the structure of physical theory regardless of what the quantity is called.

### 1.3 Why These Questions Have Not Been Asked

The holographic literature has a specific target: the Bekenstein bound on black hole entropy, its generalisation to arbitrary regions, and its implications for quantum gravity and the nature of spacetime. The goal is to understand the information-theoretic structure of gravity. Distinguishability in the operational sense — as a criterion grounded in measurement theory, specifying what a physical process can resolve — is not the natural framing for that programme.

The computational physics literature (Lloyd, Margolus, Levitin) has a different target: the maximum rate and total count of physical operations. The connection to state-space distinguishability is implicit but not developed.

The thermodynamics of information literature (Landauer, Bennett) focuses on the energetic cost of individual operations, not on the global structure of what any causal patch can support.

What has been missing is a framework that asks, explicitly and operationally: across all three literatures, what is actually being bounded, and do the three bounds form a coherent structure? The VERSF programme, by needing to use the bound in a pre-geometric context where the relevant quantity is precisely distinguishability rather than entropy, provides the motivation for asking these questions. The reason the need is acute is specific: the fold framework derives a commitment structure — a locally finite, countable partial order of irreversible distinction events — and the structural necessity of the fold boundary depends entirely on distinguishability being finite. Infinite distinguishability would allow distinctions to accumulate without bound within the void substrate, dissolving the need for a boundary at all. Finite distinguishability is therefore not an auxiliary assumption of the fold framework but the structural condition that makes the fold boundary necessary.

### 1.4 Structure of the Paper

Section 2 develops Contribution I: the operational reframing, the formal bridge from entropy to distinguishability, and what changes as a result. Section 3 develops Contribution II: the three-route convergence, its formalisation, and what it implies. Section 4 develops Contribution III: the distinguishability hierarchy, its three layers, and which layer is appropriate for which class of question. Section 5 traces the consequences of all three contributions for the Taylor's Number paper's specific predictions. Section 6 identifies open questions. Section 7 concludes.

---

## 2. Contribution I: The Operational Reframing

### 2.1 What the Holographic Literature Actually Bounds

The Bekenstein–Hawking entropy bound states that the maximum entropy of any physical system enclosed within a surface of area  $A$  is

$$S_{\max} = k_B \cdot A / (4 \ell_P^2)$$

Entropy here is a thermodynamic quantity: it counts the logarithm of the number of microstates consistent with a given macrostate. The holographic principle generalises this to the statement that the entropy of any spatial region is bounded by its surface area in Planck units. The bound is an entropy bound.

The existing literature treats the connection to information content as immediate: entropy in nats divided by  $\ln 2$  gives bits, and bits count distinguishable states. But this move — from maximum entropy to maximum distinguishable states — conceals a conceptual step that deserves to be made explicit.

Entropy counts microstates consistent with macroscopic constraints. Distinguishability counts states separable by a physical process. These coincide when every microstate is individually accessible to measurement. But the holographic literature does not prove that every microstate below the entropy ceiling is individually resolvable by a physical process. It proves that the number of microstates is bounded. Whether those microstates are all distinguishable in the operational sense is a separate question.

## 2.2 The Conceptual Gap: From Entropy to Distinguishability

To see the gap precisely, consider two scenarios:

**Scenario A.** A physical system has  $10^N$  microstates, all of which are individually distinguishable by some measurement process requiring finite resources. The entropy bound and the distinguishability count agree: both give  $N$  bits.

**Scenario B.** A physical system has  $10^N$  microstates, but quantum coherence, decoherence timescales, or thermodynamic cost constraints mean that only a subset can be individually resolved by any finite physical process. The entropy bound gives  $N$  bits; the distinguishability count gives fewer.

The holographic literature establishes the entropy bound — Scenario A's ceiling. It does not establish that every state below that ceiling is operationally distinguishable. The Bekenstein bound is an upper limit on entropy, not a count of resolvable states.

In practice, for the purposes of bounding the total distinguishable states of the causal patch, the two coincide — as we show in Section 2.3. But making the bridge explicit matters for three reasons. First, it clarifies what kind of claim the bound is: a claim about physical resolvability, not merely about thermodynamic configuration counting. Second, it connects the bound to measurement theory in a way that has direct consequences for what the bound implies about precision limits and computational capacity. Third, it provides the correct framing for the pre-geometric context of the VERSF framework, where the relevant quantity is not entropy of a thermodynamic system but distinguishability of physical states by any process.

## 2.3 The Operational Definition as Formal Bridge

We propose the following as the formal bridge between entropy bounds and distinguishability:

**Definition (Physical Distinguishability).** Two states  $\psi_1$  and  $\psi_2$  within a causal patch are *physically distinguishable* if and only if there exists a physical process  $P$  — requiring finite energy  $\Delta E < \infty$ , finite time  $\Delta t < \infty$ , and finite entropy production  $\Delta S < \infty$  — that discriminates between them with probability greater than  $\frac{1}{2} + \varepsilon$  for some  $\varepsilon > 0$  independent of the specific implementation of  $P$ .

Several features of this definition are significant.

The condition *finite energy, time, and entropy* is not merely practical. It is the condition that links distinguishability directly to the Bekenstein, Lloyd, and Landauer bounds. A process requiring infinite energy cannot be physical. A process requiring infinite time cannot complete within the causal patch. A process requiring infinite entropy production would violate the holographic bound on the total entropy of the patch. The operational definition therefore inherits its finiteness conditions directly from the physical bounds.

The probability threshold  $\frac{1}{2} + \varepsilon$ , for some  $\varepsilon > 0$  independent of implementation, ensures genuine discriminability rather than lucky guessing. It rules out distinguishability by processes that succeed only through statistical fluctuations rather than through genuine physical sensitivity to the difference.

The condition *within a causal patch* localises the definition correctly. A process that requires coordination with observers outside the causal horizon is not a physical process available to any observer within the patch.

**Conditional Bridge Proposal (Candidate formal bridge from entropy bound to distinguishability bound).** Under the operational definition above, and under the assumption that the entropy costs of independent discrimination processes are additive, the number of mutually distinguishable states within a causal patch of horizon area  $A$  is bounded above by

$$N_{\text{distinguishable}} \leq \exp( A / 4\ell_{\text{P}}^2 )$$

and therefore, applied to the cosmic horizon:

$$N_{\text{distinguishable}} \leq \mathcal{L} \approx 2.3 \times 10^{123}$$

*Proof sketch.* Suppose  $N$  mutually distinguishable states exist within the causal patch. Identifying which of  $N$  states the system is in requires at least  $\log_2(N)$  independent bits of information. By Landauer's principle, each bit of information acquired through an irreversible physical discrimination process carries a minimum entropy cost of  $k_{\text{B}} \cdot \ln 2$ . The total entropy cost of resolving  $N$  states is therefore at least  $k_{\text{B}} \cdot \log_2(N) \cdot \ln 2 = k_{\text{B}} \cdot \ln(N)$ . This total must not exceed the entropy capacity of the causal patch,  $S_{\text{max}} = k_{\text{B}} \cdot A / (4\ell_{\text{P}}^2)$ , giving  $\ln(N) \leq A / (4\ell_{\text{P}}^2)$  and therefore  $N_{\text{distinguishable}} \leq \exp(A / 4\ell_{\text{P}}^2)$ .  $\square$

This is presented as a conditional bridge proposal rather than an established result. The additivity assumption — that independent discrimination processes do not share entropy costs — is the load-bearing condition, and its status in the quantum information setting is non-trivial. Quantum entanglement between subsystems could in principle allow correlated discrimination strategies

with lower joint entropy cost than sequential independent measurements. Whether entangled strategies can exceed the sequential bound, and by how much, remains open. The argument as given shows that the entropy bound is consistent with the distinguishability bound under natural assumptions; establishing a strict equivalence requires further work. We treat the proposal as a well-motivated candidate formal bridge and flag the entangled case in Section 6 as the primary open question for Contribution I.

## 2.4 What the Reframing Changes

The operational reframing changes the character of several of the Taylor's Number paper's claims in ways that make them more precise and more defensible.

The claim that "infinity is not physically real" becomes: no physical process within the causal patch can produce or resolve more than  $\mathcal{L} \approx 2.3 \times 10^{123}$  mutually distinct outcomes. This is not a claim about mathematics. It is a precise, operational claim about the limits of physical discrimination.

The claim that divergent integrals in quantum field theory are "artifacts" becomes: the states summed over in the integral include configurations that are not operationally distinguishable within the causal patch. Integrating over them is not wrong mathematically — it is physically uninterpretable, because those configurations cannot be resolved by any physical process. The cutoffs at  $\ell_P$  and  $R_U$  mark the boundaries of operationally meaningful state space, not merely the boundaries of known physics.

The precision limit claim becomes: there exists a minimum resolvable difference between any two physical quantities, set by the requirement that the discriminating process must complete within the causal patch with finite resources. This is a claim derivable from the operational definition, not merely from the existence of a large number.

## 2.5 Connection to Measurement Theory and Finite Hilbert Space Dimension

The operational definition connects the holographic bound to the quantum theory of measurement in a way that has not been developed systematically in the cited literature. The connection runs through three fields simultaneously — quantum gravity, information theory, and measurement theory — and making it explicit is one of the central moves of this paper.

In quantum measurement theory, a measurement is described by a positive operator-valued measure (POVM) — a set of operators  $\{M_m\}$  satisfying  $\sum_m M_m^\dagger M_m = I$ , where the index  $m$  labels measurement outcomes. A measurement can distinguish states  $\psi_1$  and  $\psi_2$  if the probability distributions over outcomes differ:  $p_1(m) = \langle \psi_1 | M_m^\dagger M_m | \psi_1 \rangle \neq p_2(m) = \langle \psi_2 | M_m^\dagger M_m | \psi_2 \rangle$  for some  $m$ .

The operational definition above requires that such a POVM exists and that implementing it requires only finite physical resources. The holographic bound then constrains the Hilbert space dimension of any system within the causal patch. Specifically, a system whose Hilbert space has dimension  $D$  can be in at most  $D$  mutually orthogonal states — and orthogonal states are

maximally distinguishable. The holographic bound therefore implies that the effective Hilbert space dimension of the causal patch is bounded by

$$D_{\text{causal patch}} \leq \exp( A / 4\ell_{\text{P}}^2 ) = \exp(\mathcal{L} \cdot \ln 2)$$

This is a striking result in its own right: **the effective Hilbert space dimension accessible within the causal patch is finite**. It is not merely that we cannot access all states — it is that the total space of physically distinguishable quantum states is finite, bounded by a number derivable from the causal geometry alone. The holographic literature notes this in passing; the operational reframing makes it the central statement.

The conceptual move performed here is therefore a bridge between three domains that are usually treated separately:

```
Holographic entropy bound (quantum gravity)
  ↓ operational definition
Distinguishability limit (information theory)
  ↓ POVM constraint
Finite Hilbert space dimension (measurement theory)
```

Each arrow is a non-trivial step. The first requires Landauer's principle to connect entropy capacity to discrimination cost. The second requires the identification of orthogonal states with maximally distinguishable states under the POVM formalism. Together they establish that a bound originating in the physics of black holes has a precise and testable meaning in the language of quantum measurement: no POVM acting within the causal patch can distinguish more than  $\mathcal{L}$  mutually distinct outcomes.

## 3. Contribution II: The Three-Route Convergence

### 3.1 The Three Routes and Their Independence

Three independent theoretical frameworks arrive at  $\mathcal{L}$  as a structural limit on the distinguishability of the causal patch:

**Route 1 — Thermodynamic Optimisation (companion paper I).** Maximising an information efficiency function  $\eta(L) = B(L)/C(L)$  at quantum and cosmic scales yields two optimal length scales — near  $\ell_{\text{P}}$  and near  $R_{\text{U}}$  — whose ratio, squared holographically, gives  $\mathcal{L}$ . The starting premises are: (a) physical processes optimise information benefit per thermodynamic cost, and (b) the holographic principle determines how scale ratios translate to state counts.

**Route 2 — Empirical Constraints (companion paper II).** Four established results — the Bekenstein–Hawking bound, the holographic principle, Landauer's principle, and the Lloyd bound — independently constrain the information content of the causal patch and converge on  $\mathcal{L}$ . The starting premises are the four established results themselves, accepted independently of the VERSF programme.

**Route 3 — Pre-geometric Derivation (synthesis paper).** From Axiom V1 — that the void substrate admits no persistent informational distinctions — the fold boundary is derived as a structural necessity, and the fold's commitment structure is locally finite and countable. The finiteness of the commitment structure implies finite distinguishability of the emerging physical states, and the bound is compatible with  $\mathcal{L}$ . The starting premises are the void axiom and the topological requirements for fact formation.

**Independence of the three routes.** Route 1 does not use the Bekenstein bound, Landauer's principle, or the void axiom in its core derivation — it uses the optimisation principle and the holographic principle. Route 2 does not use the optimisation framework or the void axiom — it uses four established results whose validity predates the VERSF programme. Route 3 does not use the optimisation framework, the Bekenstein bound, or Landauer's principle — it uses only the pre-geometric axioms of the void-fold framework. Routes 1 and 2 share the holographic principle as a physical input, so their inputs are not fully disjoint; the independence claim is at the level of motivating framework and derivational structure, not at the level of all physical vocabulary used.

### 3.2 Formalising the Convergence Claim

Let  $C$  be the structural constraint: *within any causal patch, the number of mutually distinguishable states is finite and bounded above by the holographic entropy of the causal boundary.*

Let  $R_i$  denote route  $i$ , with motivating framework  $F_i$  and conclusion  $C_i$ . We claim:

1.  $F_1, F_2, F_3$  are conceptually independent at the level of derivational structure — no route is reducible to either other route, even where they share physical inputs such as the holographic principle.
2.  $C_1 = C_2 = C_3 = C$  (all three conclude the same structural constraint).

Condition 2 is the convergence. Condition 1 is the independence claim, and it requires care.

**On condition 1.** The three frameworks have distinct foundational starting points. Route 1 begins from a thermodynamic optimisation principle — the claim that physical processes maximise information benefit per unit thermodynamic cost. Route 2 begins from four empirically established results accepted independently of the VERSF programme. Route 3 begins from the pre-geometric void axiom, with no reference to thermodynamics or established physical bounds. The holographic principle appears as an input in both Routes 1 and 2, so their physical inputs are not fully disjoint. What is claimed is independence at the level of motivating framework and derivational structure: Route 1 uses the holographic principle to convert a scale ratio to a state count, having derived the scales from optimisation; Route 2 uses it as one of four independent empirical constraints; Route 3 derives finite distinguishability without invoking it at all. The three routes are not reducible to one another even where they share physical vocabulary.

**On condition 4.** All three routes conclude that physical distinguishability within the causal patch is bounded above by the holographic entropy of the causal horizon. The numerical value  $\mathcal{L} \approx 2.3$

$\times 10^{123}$  follows in each case from the same observable quantities:  $R_U$  and  $\ell_P$ . The structural claim is identical.

### 3.3 Why the Convergence Is Non-Trivial

One might object that the convergence is unsurprising because all three routes use the same physical constants and the same cosmic boundary. This objection fails for the following reason.

The space of possible structural conclusions about the causal patch is not one-dimensional. A framework could conclude that distinguishability is bounded by a different function of  $R_U$  and  $\ell_P$  — for example,  $(R_U/\ell_P)^3$  (a volume ratio) or  $(R_U/\ell_P)$  (a linear ratio) or some function involving other constants such as the fine structure constant  $\alpha$  or the proton-to-Planck mass ratio. It could conclude that distinguishability is bounded by the Lloyd count ( $\sim 10^{120}$ ) rather than the holographic ceiling ( $\sim 10^{123}$ ). It could conclude that distinguishability is not bounded at all at the causal-patch level.

The fact that all three routes — optimisation, empirical constraints, and pre-geometric derivation — converge specifically on the holographic entropy of the causal boundary, and not on any of these alternatives, is not forced by the shared use of physical constants. It reflects a common structural feature of the three frameworks: all three recognise that the holographic principle sets the geometric ceiling on state-space capacity, and all three reason toward that ceiling from different starting points.

**Why the convergence is not forced.** The convergence is not an artifact of dimensional analysis alone. Many alternative dimensionless combinations of the fundamental constants could have appeared as the natural output of one or more of the three routes: volume ratios of the form  $(R_U/\ell_P)^3$ , linear ratios  $(R_U/\ell_P)$ , combinations involving the fine structure constant  $\alpha$  or the proton-to-Planck mass ratio, or the Lloyd dynamical count  $\sim 10^{120}$  rather than the holographic ceiling  $\sim 10^{123}$ . The three routes could have converged on different combinations, or failed to converge at all. The fact that all three converge specifically on the holographic area ratio — and not on any of these alternatives — suggests that the causal boundary plays a structurally privileged role in the organisation of physical information, one that is visible from geometric, dynamical, and pre-geometric starting points independently.

### 3.4 The Convergence as Structural Evidence

The convergence has a specific evidential character. It does not prove that  $\mathcal{L}$  is the correct bound — no convergence argument can substitute for experimental confirmation. But it does provide evidence of a particular kind: evidence that the bound reflects a structural feature of physical reality that is visible from multiple theoretical perspectives, rather than an artifact of any one perspective.

This is analogous to how the convergence of thermodynamic, kinetic, and statistical mechanical arguments on the ideal gas law provides evidence that the law reflects a genuine feature of dilute gas behaviour, rather than being an artifact of any one approach. Or how the convergence of general relativity, quantum field theory in curved spacetime, and string theory on the

Bekenstein–Hawking entropy formula provides evidence that black hole entropy is a real feature of gravitational systems.

The three-route convergence on  $\mathcal{L}$  provides the same type of evidence for the causal-patch distinguishability limit: it is not a calculation artifact but a structural constraint visible across multiple independent frameworks.

The convergence also has a specific implication for the VERSF programme. If the void-fold framework (Route 3) produces a finite distinguishability constraint consistent with  $\mathcal{L}$  without knowing  $\mathcal{L}$  in advance, that is evidence that the void-fold framework captures real structure. The pre-geometric framework is not merely consistent with established physics — it independently converges on the same constraint that established physics requires.

## 4. Contribution III: The Distinguishability Hierarchy

### 4.1 The Three Bounds and Their Values

Three distinct results bound aspects of physical distinguishability within the causal patch, and they give numerically different values:

Bound	Value	Source
Holographic (Bekenstein-Bousso)	$\mathcal{L} \approx 2.3 \times 10^{123}$	Geometric: horizon area in Planck units
Lloyd	$N_{\text{ops}} \approx 10^{120}$	Dynamical: total state transitions in cosmic history
Landauer (at $T_{\text{CMB}} = 2.7 \text{ K}$ )	$N_{\text{L}} \approx 10^{90}$	Thermal: operationally accessible distinctions now

The existing literature treats these as broadly consistent — all large, all finite, all pointing in the same direction. But the gap between  $10^{123}$  and  $10^{90}$  is thirty-three orders of magnitude, which is not a rounding error. The question of whether these are three approximations to the same thing or three distinct quantities has not been systematically addressed.

### 4.2 The Hierarchy Proposal

We propose that these three bounds are not competing approximations but three layers of a single nested structure, each appropriate for a different class of physical question:

**Layer 1 — The Holographic Ceiling (geometric).** What is the maximum number of distinct states that could, in principle, be supported by the causal patch, given its geometry and the laws of quantum gravity?

**Layer 2 — The Lloyd Instantiation Limit (dynamical).** Of the states permitted by Layer 1, how many have actually been physically instantiated — brought into existence by real physical processes — over the history of the universe to date?

**Layer 3 — The Landauer Access Floor (thermal/operational).** Of the states instantiated in Layer 2, how many are currently operationally accessible — resolvable by a physical discrimination process at current thermal conditions?

The three layers are nested:  $\text{Layer 3} \leq \text{Layer 2} \leq \text{Layer 1}$ . They are not approximations to the same quantity. They answer different questions.

### 4.3 Layer 1: The Holographic Ceiling

The holographic entropy bound of the causal patch,

$$\mathcal{L} = \pi \cdot R_U^2 / \ell_P^2 \approx 2.3 \times 10^{123}$$

is a geometric quantity. It depends only on the area of the causal horizon and the Planck length. It is independent of the universe's thermal history, its age, or the current temperature of the cosmic microwave background. It would have the same value whether the universe were 13.8 billion years old or 1 billion years old or 100 billion years old, provided  $R_U$  were the same.

Layer 1 is the appropriate bound for questions about the *capacity* of the causal patch: the maximum number of states that could, in principle, be distinguished if all of the patch's information-processing resources were available. It is the answer to "how much could the universe in principle contain?"

Questions of the Layer 1 type include: What is the effective Hilbert space dimension of the causal patch? What is the maximum entropy of any physical system within the patch? What is the absolute upper bound on the information content of any physical theory that could be applied within the patch?

### 4.4 Layer 2: The Lloyd Instantiation Limit

The Lloyd bound,

$$N_{\text{ops}} \leq 2 \cdot M_U \cdot c^2 \cdot t_U / (\pi \cdot \hbar) \approx 10^{120}$$

counts the maximum number of elementary state transitions — discrete physical events — that could have occurred in the observable universe over its history from the Big Bang to now. It is a dynamical quantity, depending on the universe's mass-energy content and its age.

Layer 2 is the appropriate bound for questions about the *history* of the causal patch: the maximum number of distinct configurations that have been physically instantiated by real physical processes. A state that has never been the outcome of any physical transition — that has never been actualised — is not a physically real configuration in any operationally meaningful

sense. The Lloyd bound limits the total distinguishable configurations that physical history could have generated.

Layer 2 < Layer 1 because the universe has not used all of its geometric information capacity. The gap — approximately  $10^3$  orders of magnitude — reflects the ratio between what the geometry permits and what cosmic history has actually produced. This is not a mystery or a discrepancy; it is the expected relationship between a geometric ceiling and a dynamical history.

Questions of the Layer 2 type include: How many distinct physical configurations has the universe generated? What is the maximum complexity of any physical system that could have evolved by natural processes? What is the appropriate bound for questions about the computational history of the universe rather than its instantaneous capacity?

### 4.5 Layer 3: The Landauer Access Floor

The Landauer bound at current CMB temperature  $T_{\text{CMB}} \approx 2.7 \text{ K}$ ,

$$N_{\text{L}} \leq F / (k_{\text{B}} \cdot T_{\text{CMB}} \cdot \ln 2)$$

where  $F$  is the available free energy, gives approximately  $10^{90}$  for conservative estimates of accessible free energy at current cosmic conditions. This is a thermal quantity, depending on the current temperature of the universe and the available free energy for information processing.

Layer 3 is the appropriate bound for questions about *current operational access*: the maximum number of distinctions that can be thermodynamically registered by any physical process operating at present conditions. A distinction that would cost more thermodynamic free energy to register than is currently available in the accessible environment cannot be operationally realised now, even if it exists as a physical configuration (Layer 2) and is permitted by the geometry (Layer 1).

Layer 3 < Layer 2 because the universe's increasing entropy and declining available free energy — as it approaches thermal equilibrium — reduces the thermodynamically accessible fraction of the configurations that have been physically instantiated. This is also not a mystery; it reflects the thermodynamic arrow of time and the increasing entropic cost of fine-grained discrimination as the universe cools and dilutes.

Questions of the Layer 3 type include: What is the current precision limit of any measurement? What is the maximum number of distinct computational states accessible to any physically realisable computer operating today? What is the thermodynamically relevant bound for current experimental predictions?

### 4.6 The Nested Structure and Its Implications

The nested structure  $\text{Layer 3} \leq \text{Layer 2} \leq \text{Layer 1}$  is not merely a notational convenience. It reflects a genuine physical distinction between three types of limitation on distinguishability:

- **Geometric limitations** (Layer 1) are set by the laws of quantum gravity and the causal structure of spacetime. They are the hardest constraints — no physical process within the patch can exceed them.
- **Dynamical limitations** (Layer 2) are set by the actual physical history of the universe. They reflect what physical processes have generated, not what the geometry permits. They can in principle be approached from below (by maximising the use of available energy for physical transitions) but never exceeded.
- **Thermal/operational limitations** (Layer 3) are set by current thermodynamic conditions. They are the softest constraints in the sense that they change over time as the universe evolves — they were less restrictive in the early universe when temperatures were high and free energy abundant, and they become more restrictive as the universe cools.

**Implication for physical claims.** A claim about the maximum capacity of the causal patch should use Layer 1. A claim about the maximum complexity of physical systems that could have naturally evolved should use Layer 2. A claim about current measurement precision limits should use Layer 3.

The Taylor's Number paper's claims span all three layers:

- The claim that divergences in QFT are cutoff by the existence of a minimum and maximum physically meaningful scale is a Layer 1 claim — it concerns the geometric structure of the causal patch regardless of its thermal history.
- The claim that the total number of meaningful computational steps is bounded is a Layer 2 claim — it concerns the dynamical history of physical processes.
- The claim that measurement precision will saturate as technology improves is a Layer 3 claim — it concerns the thermodynamically accessible resolution at current and near-future conditions.

Distinguishing which layer each claim belongs to sharpens the predictions and clarifies what evidence would confirm or falsify each.

---

## 5. Consequences: What Each Contribution Implies

### 5.1 Consequences of the Operational Reframing

**For the QFT regularisation argument.** The Taylor's Number paper argues that divergent integrals in QFT are physically meaningless because they sum over states outside the physically relevant domain. The operational reframing makes this precise: the states summed over beyond  $\ell_{\text{P}}^{-1}$  and  $R_{\text{U}}^{-1}$  are not operationally distinguishable — no physical process within the causal patch, with finite resources, can discriminate between configurations that differ only at sub-Planck scales or across super-horizon distances. The cutoffs are not arbitrary regularisation choices. They mark the boundaries of the operationally meaningful Hilbert space of the causal patch.

**For the precision measurement limit.** The prediction that no measurement can achieve relative precision better than  $1/\mathcal{L} \approx 4 \times 10^{-124}$  is now a Layer 1 claim: an absolute geometric ceiling on discriminability. But the more immediately testable prediction is a Layer 3 claim: precision limits will saturate at the Landauer access floor as technology improves and the thermodynamic cost of finer discrimination becomes inaccessible. These are distinct predictions with distinct timescales and distinct experimental signatures.

**For the halting argument.** The claim that physical computation must eventually halt because it cannot exceed  $\mathcal{L}$  state distinctions is properly a Layer 2 claim: the dynamical history of the universe can generate at most  $\sim 10^{120}$  distinct transitions, not  $10^{123}$ . The halting argument should reference the Lloyd instantiation limit, not the holographic ceiling, for claims about the computational history of physical systems.

## 5.2 Consequences of the Convergence Result

The three-route convergence implies that the finite distinguishability constraint is visible from three independent perspectives. This has a specific consequence for the falsifiability of the VERSF programme: any experimental result that violated the constraint — any observation of more than  $\mathcal{L}$  distinguishable states in a causal patch — would simultaneously falsify all three routes. No rescue by adjusting one framework while preserving the others would be available.

This is a strong property. Most theoretical frameworks can be locally modified to accommodate anomalous observations. A constraint supported by three mutually independent routes is more resistant to local modification, because any modification to one route leaves two others intact to reassert the constraint. The convergence therefore provides not just evidential support for the constraint but structural robustness.

## 5.3 Consequences of the Hierarchy

The hierarchy proposal resolves a potential confusion in the Taylor's Number paper's treatment of the bounds. The paper uses the holographic bound  $\mathcal{L}$  as the primary quantity throughout, but some of its predictions are more naturally Layer 2 or Layer 3 claims. Mapping each prediction to its appropriate layer produces a more precise and falsifiable set of predictions:

Prediction	Appropriate Layer	Relevant Bound
QFT divergence cutoffs	Layer 1 (geometric)	$\mathcal{L} \approx 10^{123}$
Maximum physically evolved complexity	Layer 2 (dynamical)	$N_{\text{ops}} \approx 10^{120}$
Current precision measurement limit	Layer 3 (thermal)	$N_L \approx 10^{90}$
408-qubit quantum computing prediction	Layer 1 or 2?	Unclear — see §6
Eventual precision saturation	Layer 3 $\rightarrow$ Layer 1 asymptote	Both

The mapping from a causal-patch distinguishability bound to laboratory-scale quantum systems remains an open problem. The 408-qubit figure derives from  $\log_2(\mathcal{L}) \approx 410$ , making it a Layer 1 claim about the geometric ceiling of the entire causal patch. But a laboratory quantum computer occupies a negligible fraction of the causal patch's volume, and the geometric ceiling of a

laboratory-scale region is vastly larger than 408 qubits. A mechanism connecting the global bound to local system behaviour — whether through a quantum gravity effect, a decoherence argument, or a structural property of the fold commitment structure — is needed before this prediction can be placed in the hierarchy with confidence.

---

## 6. Open Questions

**The Bridge Proposition and entangled discrimination strategies.** The Bridge Proposition in Section 2.3 establishes the connection between the Bekenstein entropy ceiling and the operational distinguishability count under the assumption that independent discrimination processes have additive entropy costs. Whether quantum entanglement between subsystems allows correlated discrimination strategies with lower joint entropy cost — and whether such strategies could in principle allow more than  $\exp(A / 4\ell_P^2)$  states to be distinguished within a causal patch — is a question in quantum information theory that the proposition does not resolve. A full proof of the bridge would require showing either that entangled strategies cannot beat the sequential bound, or that they can but only up to a bounded correction factor that does not change the order of magnitude of the distinguishability limit.

**The laboratory propagation mechanism.** All three contributions sharpen the theoretical structure of the causal-patch distinguishability limit, but none resolves the problem of how a causal-patch bound propagates to laboratory scales. The 408-qubit prediction requires a mechanism connecting  $\mathcal{L}$  to the behaviour of a system occupying  $\sim 10^{-80}$  of the causal patch's volume. Identifying that mechanism — whether through a quantum gravity effect, a decoherence argument, or a structural property of the fold commitment structure — is the most pressing open problem for the programme's experimental predictions.

**The hierarchy across cosmic time.** Layer 3 (the Landauer access floor) is time-dependent: it was higher in the early universe when free energy was abundant and lower now as the universe cools. A complete treatment of the hierarchy should include its time evolution — how the three layers relate to each other at different epochs, and whether the approach to thermal equilibrium produces any observable signatures as Layer 3 approaches Layer 2 from below over cosmological timescales.

**Convergence in other physical contexts.** The three-route convergence established in Section 3 applies to the causal patch of the observable universe. Does an analogous convergence hold for other regions — for example, for a black hole horizon, for a Rindler wedge in flat spacetime, or for a de Sitter static patch? If the convergence is a general feature of any causal horizon rather than specific to the cosmological case, that would significantly strengthen the structural interpretation of the result.

---

## 7. Conclusion

The causal-patch holographic distinguishability limit  $\mathcal{L} \approx 2.3 \times 10^{123}$  is not a new number. It is the holographic entropy bound of the observable universe, known in the literature for decades. The original theoretical work of the VERSF programme concerning this quantity consists of three structural contributions, developed in this paper.

**Contribution I** establishes the operational reframing: the formal bridge from the holographic entropy bound to the count of operationally distinguishable states, using the operational definition of physical distinguishability as the connector. This move connects the holographic bound to measurement theory and provides a precise criterion for what it means for a physical difference to be real rather than merely mathematical. The bridge goes through Landauer's principle, which provides the minimum entropy cost of any discrimination event and thereby links the thermodynamic entropy ceiling to the operational distinguishability count.

**Contribution II** formalises the three-route convergence: the observation that a thermodynamic optimisation, four empirical constraints, and a pre-geometric void-fold derivation independently arrive at the same structural constraint, drawing on distinct motivating frameworks with partially overlapping physical inputs. The convergence is shown to be non-trivial and is interpreted as structural evidence that finite distinguishability is a feature of physical reality visible across multiple independent theoretical perspectives, not an artifact of any single framework.

**Contribution III** proposes the distinguishability hierarchy: the reorganisation of the holographic, Lloyd, and Landauer bounds as three layers of a single nested structure — geometric ceiling, dynamical instantiation limit, and operational access floor — each appropriate for a different class of physical claim. The hierarchy resolves the apparent inconsistency between the three bounds' numerical values and provides a principled basis for selecting the appropriate bound for each of the Taylor's Number paper's specific predictions.

Together, these three contributions constitute the VERSF programme's original theoretical work on the causal-patch distinguishability limit. They do not depend on the claim that  $\mathcal{L}$  is a new constant. They depend on developing systematically, as a unified framework, structural questions about a known quantity that the existing literature has not treated in this combined form.

---

## References

### Holographic principle and entropy bounds

- Bekenstein, J.D. (1972). Black holes and the second law. *Lettere al Nuovo Cimento* 4, 737.
- Bekenstein, J.D. (1973). Black holes and entropy. *Physical Review D* 7, 2333.
- Hawking, S.W. (1975). Particle creation by black holes. *Communications in Mathematical Physics* 43, 199.
- 't Hooft, G. (1993). Dimensional reduction in quantum gravity. *arXiv:gr-qc/9310026*.
- Susskind, L. (1995). The world as a hologram. *Journal of Mathematical Physics* 36, 6377.

- Bousso, R. (1999). A covariant entropy conjecture. *Journal of High Energy Physics* 9907, 004.
- Bousso, R. (2002). The holographic principle. *Reviews of Modern Physics* 74, 825.

### **Computational and thermodynamic bounds**

- Landauer, R. (1961). Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development* 5, 183.
- Lloyd, S. (2000). Ultimate physical limits to computation. *Nature* 406, 1047.
- Lloyd, S. (2002). Computational capacity of the universe. *Physical Review Letters* 88, 237901.
- Margolus, N., Levitin, L.B. (1998). The maximum speed of dynamical evolution. *Physica D* 120, 188.
- Bérut, A., et al. (2012). Experimental verification of Landauer's principle. *Nature* 483, 187.
- Bennett, C.H. (1982). The thermodynamics of computation — a review. *International Journal of Theoretical Physics* 21, 905.

### **Quantum measurement theory**

- Helstrom, C.W. (1976). *Quantum Detection and Estimation Theory*. Academic Press.
- Holevo, A.S. (1982). *Probabilistic and Statistical Aspects of Quantum Theory*. North-Holland.
- Nielsen, M.A., Chuang, I.L. (2000). *Quantum Computation and Quantum Information*. Cambridge University Press.