

Toward a Principle of Significance

*A Holographic Bound on the Physically Distinguishable Configuration Space of a Bounded Observer
(Once more on the finitude of the Universe)*

Sergej Materov
Independent Researcher
© September 2026

Abstract

The principle of causality bounds what can influence a given point: nothing outside its past light cone. This note asks whether a complementary bound exists on what can be physically significant to a bounded observer - i.e. on the number of configurations that observer could, even in principle, operationally distinguish. We show that such a bound follows from Bousso's covariant entropy bound, established in the relevant semiclassical setting under stated assumptions, together with elementary one-shot quantum information theory, with no Banks-Fischler-type postulate anywhere in the derivation. Once the notion of operational significance (Definition 3.2) is sharpened into an explicit N-ary ensemble-discrimination task (Definition 3.2'), and a physical “configuration” is represented as a classical register tensored with the light-sheet Hilbert space, §4.6 proves

$$N(O) \leq \exp \left\{ \frac{\left[\frac{A_{\max}(O)}{4\ell_p^2} + h(\varepsilon) \right]}{1 - \varepsilon} \right\}$$

directly from Bousso's theorem, the data-processing inequality, and the Holevo bound — no further conjecture is invoked. This bound is proved for a single ε -significant N-ary ensemble in the precise sense of Definition 3.2'; whether it also bounds every configuration ever operationally significant to O across its entire history, in the broader sense of Definition 3.2 that motivates the principle informally, is item (i)(c) of §7 and is not itself established here. The zero-error bound $N(O) \leq \exp \left[\frac{A_{\max}(O)}{4\ell_p^2} \right]$ would formally correspond to the $\varepsilon \rightarrow 0$ limit, though taking that limit rigorously - recovering an exact, error-free statement - is not established here (item (i)(a) of §7). §3.3 and §4.7 separately address the finiteness of $A_{\max}(O)$ for FRW observers whose apparent-horizon area stays bounded over their history ($\Lambda > 0$, closed recollapsing cosmologies, or finite histories; for open or flat $\Lambda = 0$ cosmologies with unbounded history the global quantity is infinite and only a stage-local version is finite), and describe it as the supremum of a monotone quantity along a holographic screen for fully dynamical horizons (via the Bousso–Engelhardt area law), so that the stage-local bound applies beyond static horizons and, at the stage-local level, without restriction to $\Lambda > 0$. We note the independent numerical cross-check offered by Lloyd's (2002) computational-capacity bound - obtained by an entirely unrelated argument and landing on the same order of magnitude, $\sim 10^{120}$ - and list what would be needed to upgrade the conjectural component to a proof.

1. Motivation: Causality and a Missing Complement

The causal structure of a Lorentzian spacetime already gives physics one half of a natural pair of bounds. The light cone answers the question “what can possibly affect this point?” with a sharp, geometric, and rigorously established negative answer: nothing outside $J^-(p)$, the causal past of p . No comparably sharp answer is standardly given to the complementary question: “how much can this point - or, more precisely, a physical observer whose worldline passes through it - ever come to distinguish, and hence regard as physically significant?” The causality principle bounds propagation; what might be called a principle of significance would bound capacity.

This is not merely a rhetorical parallel. The physics needed to give the second question a rigorous partial answer already exists, developed for a different purpose - the thermodynamics of horizons - over the past five decades. This note assembles that physics into an explicit, conditionally stated principle, and is careful throughout to mark the line between what is proved and what is proposed.

1.1 Related work. The building blocks used in §4.6 are individually established, and the closest prior combinations of them are closer than a first look suggests. Bousso (2017) already combines the quantum Bousso bound directly with the Holevo theorem,

exactly the architecture of §4.6 - so that combination itself is not the novelty claimed here. Bousso's bound, however, is a bound on communication channel capacity between two distant systems (χ of order $E\Delta t/\hbar$), not on the number of configurations distinguishable by a bounded observer, and it does not include an analogue of (S2')'s explicit light-sheet-to-observer channel. Bao & Ooguri (2017) apply the Holevo information to a different, AdS/CFT-specific question - distinguishability of black hole microstates by measurements on a CFT subregion - not a general bounded-observer framework. Van Acoleyen (2010) derives a local operational form of the covariant bound - the number of bits decodable on a light-sheet interval - from an uncertainty relation on the measuring screen, by the author's own account a heuristic argument whose optimality is not established; Chronis & Sifakis (2025) develop cosmological consequences of a causal patch's operational Hilbert-space dimension, taken as a conditional starting hypothesis rather than derived. To our knowledge, none of these gives an explicit one-shot formulation connecting the covariant entropy bound, via an entropy-non-increasing physical channel from a light-sheet to an observer-accessible subsystem (S2'), to a finite-error bound $N\epsilon$ on the number of configurations that subsystem can discriminate (Theorem 4.6.4); that specific combination, not the covariant bound, the Holevo bound, or their pairing individually, is this paper's claimed contribution.

2. Background: From Bekenstein to Bousso

2.1 The Bekenstein bound. Bekenstein (1973, 1981) established that the entropy of any bounded physical system of energy E confined to a sphere of radius R satisfies $S \leq \frac{2\pi RE}{\hbar c}$ is a universal bound relating entropy to size and energy, independent of the system's internal composition.

2.2 Black hole area law. Bekenstein's proposal that black holes carry entropy proportional to horizon area was given a precise coefficient by Hawking's discovery of black hole radiation: $S_{\text{BH}} = \frac{A}{4\ell_p^2}$ (in units where $k_B = c = \hbar = G = 1$), the Bekenstein–Hawking entropy.

2.3 Cosmological horizons. Gibbons and Hawking (1977) extended this area law from black hole horizons to cosmological event horizons: an observer in a spacetime with positive cosmological constant is surrounded by a horizon of finite area, radiating at a temperature set by its surface gravity and carrying entropy $S = \frac{A}{4\ell_p^2}$.

2.4 The covariant entropy bound. Bousso (1999) formulated, as a conjecture, a generalization of the area-law intuition not tied to horizons specifically; it has since been established in several semiclassical regimes and under stated hypotheses (see the discussion of proof regimes in §4). Given any surface of area A and either family of orthogonal null geodesics whose expansion is non-positive (a “light-sheet”), the entropy of matter crossing that light-sheet satisfies $S \leq \frac{A}{4\ell_p^2}$, where the classical arguments rest on the focusing theorem (Raychaudhuri's equation) under the null energy condition together with further hypotheses on the entropy current (see §4). This is the core of everything that follows, used here under those stated hypotheses.

2.5 The Banks–Fischler proposal. Banks (2000) and Fischler (2000) proposed elevating the finiteness of cosmological horizon entropy to a defining principle: that the true, complete Hilbert space describing a de Sitter universe is finite-dimensional, of dimension $\exp(S_{\text{dS}})$. This remains an actively contested proposal, not a theorem, and is noted in the literature to require an additional, non-automatic assumption connecting bounded entropy to finite Hilbert-space dimension.

3. Formal Setup

Definition 3.1 (Bounded observer). A bounded observer O is a timelike worldline together with a causally bounded patch $D(O)$ - a causal diamond or a static patch bounded by a horizon - such that the boundary $\partial D(O)$ has finite area at every cross-section reachable along O 's worldline. Write $A_{\text{max}}(O)$ for the supremum of this area over O 's entire history.

Definition 3.2 (Operational significance). A physical configuration is significant to O if there exists, in principle, a measurement protocol confined to $D(O)$ - using only resources and time available within the patch over O 's history - that can distinguish it from at least one alternative configuration with non-vanishing probability. This is an operational, not merely ontological, notion (compare Remark 3.4(ii)'s operational/ontological distinction in the companion No-Go paper, Materov 2026a). A general, theory-independent physical criterion for what makes a bounded subsystem an observer at all - prior to and independent of this paper's causal-patch framework - is developed separately in Materov (2026d); the tensor-factor unification of that criterion with the operational representation of §4.6.0 above remains open.

3.3 Finiteness of $A_{\max}(O)$ for FRW Spacetimes

Definition 3.1 presupposes $A_{\max}(O)$ is finite. This is automatic for an observer inside a de Sitter horizon ($\Lambda > 0$) or inside a black hole exterior with a finite apparent horizon; a spatially flat or open FRW universe with $\Lambda = 0$, however, has no finite cosmological event horizon at all.

The resolution is to read $\partial D(O)$ in Definition 3.1 as the apparent (Hubble) horizon rather than the event horizon whenever the two diverge. For any FRW cosmology, expanding or not, and regardless of the sign of Λ , the apparent horizon - the outermost marginally trapped surface, at areal radius $r_A = \frac{1}{H}$ for a spatially flat FRW metric - has finite area at every instant of O 's history at which H (more generally, the effective density $8\pi\rho + \Lambda$) is non-zero, which holds at every finite time of an expanding cosmology away from a genuine curvature singularity but excludes, e.g., the turning point of a flat recollapsing history. Cai & Kim (2005), building on the trapping-horizon thermodynamics of Bak & Rey, show that the first law of thermodynamics $dE = TdS - PdV$ holds on the apparent horizon with exactly the Bekenstein-Hawking form $S = \frac{A}{4\ell_P^2}$ for FRW cosmologies governed by Einstein gravity, independent of Λ . Bousso's light-sheet construction (§2.4) itself needs only a surface of non-positive expansion, which the apparent horizon supplies by definition - it does not require that the surface in question be an event horizon in the causal, teleological sense.

With this reading, $A_{\max}(O) := \sup$ over O 's history of the apparent-horizon area is finite exactly when that area stays bounded along the history. For any spatial curvature the Friedmann equation gives $H^2 + k/a^2 = (8\pi/3)\rho + \Lambda/3$ (Planck units), so the apparent-horizon radius is $r_A = [(8\pi/3)\rho + \Lambda/3]^{-1/2}$ and its area is $A = 12\pi/(8\pi\rho + \Lambda)$. Hence $A_{\max}(O) < \infty$ if and only if $\inf(8\pi\rho + \Lambda) > 0$ along O 's history. This holds for $\Lambda > 0$ ($A \leq 12\pi/\Lambda$, given $\rho \geq 0$), for closed recollapsing cosmologies with $\Lambda = 0$ (ρ is bounded below by its value at maximum expansion), and for any observer with a finite history on which $8\pi\rho + \Lambda$ stays positive. It fails for open or spatially flat $\Lambda = 0$ cosmologies that expand forever, where $\rho \rightarrow 0$: for flat matter domination $A = 9\pi^2$ and for flat radiation domination $A = 16\pi^2$, so $A_{\max}(O) = \infty$ for an observer with unbounded history. In those cases (S1) is not satisfied and Theorem 4.3 does not apply (its right-hand side would in any case be $\exp(\infty)$); (S1) therefore remains an assumption in general, discharged in the cases just listed, and still needs $\Lambda > 0$ or one of the other conditions as an input.

What survives without such a restriction is a stage-local quantity. Define $A_{\max}(O;t)$ as the supremum of the apparent-horizon area over O 's history up to stage t . It is finite whenever $8\pi\rho + \Lambda$ stays positive on that interval - in particular at every finite time of an expanding flat or open FRW cosmology, whether matter-, radiation- or Λ -dominated. Theorem 4.6.4 is a stage-local statement, and its proof uses $A_{\max}(O)$ only as an upper bound on the area of the cross-section whose light-sheet carries the matter of (S2'); it therefore holds with $A_{\max}(O)$ replaced by $A_{\max}(O;t)$ (or, where the relevant light-sheet is that of the stage- t cross-section, by $A(t)$ itself), giving a non-trivial bound at every finite stage of an observer in a $\Lambda = 0$ open or flat FRW cosmology. The price is the one already stated in §4.6.5: this bounds $N_e(O;t)$ at stage t , not the historical union, and the all-history Theorem 4.3 is available only where $A_{\max}(O)$ is finite.

What is lost is a small amount of interpretive cleanliness: the apparent horizon, unlike the event horizon, is an instantaneous, observer-frame-dependent construction rather than a globally defined causal boundary, so Theorem 4.3 under these reading bounds significance relative to the patch enclosed by the current apparent horizon at each instant, not relative to a single fixed causal patch accumulated over all of O 's history. This is a genuine narrowing of the claim, stated here rather than left implicit. The open or flat $\Lambda = 0$ case with unbounded history is left open at the level of the global bound (§7(ii)).

4. The Principle of Significance

Assumptions. (S1) O is a bounded observer per Definition 3.1 (as clarified in §3.3), with $A_{\max}(O)$ finite. (S2) The null energy condition holds for all matter crossing the light-sheets of $\partial D(O)$. (S2') Light-sheet realization: the operational system B of §4.6.0 is obtained from the degrees of freedom crossing the relevant light-sheet, in the sense that there is a state ρ_{LS} of that matter and a physical process — the observer's measurement apparatus together with any intervening dynamics — taking it to ρ_B , satisfying $S(\rho_B) \leq S(\rho_{LS})$. (S3b) [Established in §4.6, conditional on (S2')] The number of operationally distinguishable configurations significant to O , in the sense of Definition 3.2, accessible over O 's entire history, is bounded by the exponential of the maximal geometric entropy the covariant bound permits on $\partial D(O)$. (S3a), the stronger, purely ontological Banks–Fischler claim, is discussed in §4.5 and is sufficient but not necessary for (S3b).

Theorem 4.3 (Principle of Significance, conditional). Under (S1), (S2) and (S3b), the number $N(O)$ of configurations ever significant to a bounded observer O satisfies

$$N(O) \leq \exp \left[\frac{A_{\max}(O)}{4\ell_P^2} \right].$$

Proof. By (S2) and Bousso's theorem, the entropy crossing any light-sheet of any cross-section of $\partial D(O)$ is bounded by (area of that cross-section)/ $4\ell_P^2$; taking the supremum over O 's history and applying (S1), the entropy bound on the light-sheets of the maximal cross-section is $\frac{A_{\max}(O)}{4\ell_P^2}$. By (S3b), this geometric entropy bound directly upper-bounds the number of operationally distinguishable configurations significant to O , giving $N(O) \leq \exp\left[\frac{A_{\max}(O)}{4\ell_P^2}\right]$ directly. ■

Status. The inference from (S1)+(S2) to the light-sheet entropy bound is the covariant entropy bound, used here under the semiclassical assumptions stated below. Bousso's original (1999) formulation was a conjecture; Flanagan, Marolf & Wald (1999) subsequently derived it from the null energy condition together with a phenomenological hypothesis on the entropy current, and Bousso, Casini, Fisher & Maldacena (2014) proved a quantum version for free fields in the weak-gravitational-backreaction limit without assuming the null energy condition at all. This paper does not adjudicate which proof regime applies to the matter relevant to a given O ; (S2)'s null energy condition is the working hypothesis used throughout. The bound on $N(O)$ via (S3b) is derived, not conjectured, in §4.6 - but only for the single ε -significant N -ary ensemble of Definition 3.2' (Theorem 4.6.4), not yet for “every configuration ever significant to O ” in the fuller sense Theorem 4.3 is stated in. Closing that gap is open problem (i)(c) of §7. Theorem 4.3 should therefore be read and cited as conditional on that gap for its historical (all-of- O 's-history) reading:

“If the geometric entropy bound on a causal horizon also bounds the operational distinguishability capacity of an observer confined within it across its entire history - and not merely at a single measurement stage - THEN $N(O)$ is bounded as stated.”

4.5 Clarifying (S3): Two Non-Equivalent Claims

(S3a) Ontological / dimension claim (literal Banks–Fischler). The Hilbert space describing the physics accessible within $D(O)$ is finite-dimensional, with $\dim = \exp(S)$. **(S3b) Operational / distinguishability claim (what Theorem 4.3 needs).** A resource-bounded observer confined to $D(O)$ cannot operationally distinguish more than $\exp(S)$ configurations, in the sense of Definition 3.2.

Lemma 4.5.1 (One-directional implication). $(S3a) \Rightarrow (S3b)$, but not conversely.

Proof. If $\dim H_{D(O)} = \exp(S)$ (S3a), then no more than $\exp(S)$ mutually orthogonal states exist under any measurement whatsoever, resource-unrestricted included, so $N(O) \leq \exp(S)$ follows immediately, giving (S3b). The converse fails: (S3b) is compatible with $\dim H_{D(O)}$ being infinite - a Holevo-type bound can limit operational distinguishability to $\exp(S)$ or less even when the underlying state space is unbounded, because time- and energy-limited measurements access only a bounded amount of information regardless of the size of the space of possible states. ◀

This relocates the open problem from an ontological claim about Hilbert-space dimension (S3a) to an information-theoretic claim about entanglement entropy on a light-sheet, via the Quantum Focusing Conjecture (Bousso, Fisher, Leichenauer & Wall, 2016; proved as the QNEC by Bousso, Fisher, Koeller, Leichenauer & Wall, 2016, and given a fully general proof via modular theory by Balakrishnan, Faulkner, Khandker & Wang, 2019). §4.6 below closes (S3b) directly.

4.6 Closing (S3b): A One-Shot Derivation from Bousso's Theorem and a Light-Sheet Realization Assumption

This section derives (S3b) directly, rather than merely situating it near established one-shot quantum-information territory. It requires one precisification of Definition 3.2 and one physical representation assumption, both stated explicitly; given those, the resulting bound on $N(O)$ follows by proof, not by a further named conjecture.

4.6.0 Operational Representation Assumption

The derivation below does not require the existence of a globally preferred tensor-product decomposition of the complete quantum-gravitational Hilbert space. It requires only an operational representation of the degrees of freedom accessible to the observer during the discrimination experiment.

Specifically, for the purpose of an N -ary discrimination task, assume that there exists a quantum system B and a classical label register X such that the candidate configurations relevant to the task are represented by the classical-quantum state $\frac{1}{N} \sum_i |i\rangle\langle i|_X \otimes \rho_i^B$, where B contains precisely the degrees of freedom on which the observer's admissible measurement acts.

The assumption is therefore not that the complete physical Hilbert space factorizes globally as $H_X \otimes H_B$, but that the operational discrimination problem admits this classical-quantum representation. The role of B in the proof is correspondingly restricted: the decoder is a POVM on B , and the averaged state $\bar{\rho}^B = \frac{1}{N} \sum_i \rho_i^B$ is the single state to which the covariant entropy bound is subsequently applied. This last step is not automatic from the operational representation alone: §4.6.0 characterizes B only as the

degrees of freedom the observer's measurement acts on, which need not coincide with the matter crossing $\partial D(O)$'s light-sheet. Applying the covariant bound to $\bar{\rho}B$ requires the further assumption (S2') below, that B is obtained from the light-sheet-crossing degrees of freedom by a process that does not increase entropy.

If no such operational representation exists for a proposed class of configurations, Theorem 4.6.4 does not apply to that class. The result is therefore conditional on an operational representation assumption, rather than on the stronger ontological assertion that the observer's accessible region possesses an independently defined tensor-product Hilbert-space factor.

Definition 3.2' (N-ary ensemble discrimination). Let t denote an admissible measurement stage along the observer's worldline. Let $\{\rho_1, \dots, \rho_N\}$ be N candidate configurations represented operationally as states on the observer-accessible system B , with classical labels recorded in X as specified in §4.6.0. Say that the ensemble is ε -significant to O at stage t if there exists a single POVM $M_1, \dots, M_N$ on B , with $\sum_i M_i \leq I$, such that under a uniform prior the average success probability satisfies $\frac{1}{N} \sum_i \text{Tr}(M_i \rho_i) \geq 1 - \varepsilon$.

This definition is deliberately stage-local. It formalizes joint operational distinguishability by one measurement protocol at one admissible stage, rather than the union of all configurations that may become distinguishable at different stages of the observer's history.

Lemma 4.6.1 (One-shot converse: hypothesis-testing entropy vs. relative entropy). For any states ρ, σ and any $\varepsilon \in (0, 1/2)$, the hypothesis-testing relative entropy $D_H^\varepsilon(\rho \parallel \sigma) := -\ln \min \{ \text{Tr}[Q\sigma] : 0 \leq Q \leq I, \text{Tr}[Q\rho] \geq 1 - \varepsilon \}$ (Wang & Renner, 2012) satisfies $D_H^\varepsilon(\rho \parallel \sigma) \leq \frac{[D(\rho \parallel \sigma) + h(\varepsilon)]}{1 - \varepsilon}$, where D is the ordinary quantum relative entropy and $h(\varepsilon) = -\varepsilon \ln \varepsilon - (1 - \varepsilon) \ln(1 - \varepsilon)$.

Proof. Let Q^* achieve the minimum defining $D_H^\varepsilon(\rho \parallel \sigma)$, so $\text{Tr}[Q^* \rho] \geq 1 - \varepsilon$. The binary measurement $\{Q^*, I - Q^*\}$ maps ρ, σ to classical distributions $p = (p_1, p_2)$, $q = (q_1, q_2)$ with $p_1 = \text{Tr}[Q^* \rho] \geq 1 - \varepsilon$, $q_1 = \text{Tr}[Q^* \sigma]$. By the data-processing inequality for relative entropy under measurement channels, $D(\rho \parallel \sigma) \geq D(p \parallel q) = p_1 \ln(p_1/q_1) + p_2 \ln(p_2/q_2)$. Since $q_2 = 1 - q_1 \leq 1$, $\ln(p_2/q_2) \geq \ln p_2$, so $D(p \parallel q) \geq p_1 \ln(p_1/q_1) + p_2 \ln p_2 = -H(p) - p_1 \ln q_1$, where $H(p)$ is the (natural-log) binary entropy of p . Hence $-\ln q_1 \leq [D(\rho \parallel \sigma) + H(p)]/p_1$. For $\varepsilon < 1/2$, $p_1 \geq 1 - \varepsilon > 1/2$ puts p on the decreasing branch of H , so $H(p) \leq H(1 - \varepsilon) = h(\varepsilon)$; together with $p_1 \geq 1 - \varepsilon$ this gives $-\ln q_1 \leq [D(\rho \parallel \sigma) + h(\varepsilon)]/(1 - \varepsilon)$. Since $q_1 = \text{Tr}[Q^* \sigma]$ equals the defining minimum, $D_H^\varepsilon(\rho \parallel \sigma) = -\ln q_1$, giving the claim. $\angle$

Lemma 4.6.2 (Ensemble cardinality via hypothesis testing). If $\{\rho_1, \dots, \rho_N\}$ is ε -significant to O in the sense of Definition 3.2', then for every reference state σ^B on H_B , $\ln N \leq D_H^\varepsilon(\rho^{XB} \parallel \pi^X \otimes \sigma^B)$, where $\rho^{XB} := \frac{1}{N} \sum_i |i\rangle\langle i|^X \otimes \rho_i^B$ is the classical-quantum extension of the ensemble and π^X is the uniform distribution on $\{1, \dots, N\}$.

Proof. Let $\{M_i\}$ be the decoder achieving ε -significance and set $Q := \sum_i |i\rangle\langle i|^X \otimes M_i$, a valid operator with $0 \leq Q \leq I$ since $\sum_i M_i \leq I$. Then $\text{Tr}[Q \rho^{XB}] = \frac{1}{N} \sum_i \text{Tr}[M_i \rho_i] \geq 1 - \varepsilon$, so Q is an admissible test at type-I error ε . Its type-II error is $\text{Tr}[Q(\pi^X \otimes \sigma^B)] = \frac{1}{N} \sum_i \text{Tr}[M_i \sigma^B] \leq \frac{1}{N} \text{Tr}[(\sum_i M_i) \sigma^B] \leq \frac{1}{N}$. This particular Q is one admissible test, so the true minimum type-II error over all admissible tests is itself $\leq \frac{1}{N}$. Since D_H^ε is defined as $-\ln[\text{min type-II error}]$ and $-\ln$ is decreasing, $D_H^\varepsilon(\rho^{XB} \parallel \pi^X \otimes \sigma^B) = -\ln[\text{min type-II error}] \geq -\ln\left(\frac{1}{N}\right) = \ln N$. $\angle$

Lemma 4.6.3 (Holevo minimization). $\min_{\sigma^B} D(\rho^{XB} \parallel \pi^X \otimes \sigma^B) = \chi := S(\bar{\rho}^B) - \frac{1}{N} \sum_i S(\rho_i)$, the Holevo quantity of the ensemble $\{\frac{1}{N}, \rho_i\}$, with $\bar{\rho}^B := \frac{1}{N} \sum_i \rho_i$ and the minimum attained at $\sigma^B = \bar{\rho}^B$. This is the standard fact underlying the Holevo bound (Nielsen & Chuang, 2000, §12.1.1) and is not reproved here. Since $S(\rho_i) \geq 0$ for every i , $\chi \leq S(\bar{\rho}^B)$.

Theorem 4.6.4 (Principle of Significance from Bousso's theorem and (S2'), fixed ε). Under (S1), (S2), (S2'), and Definition 3.2', for any $\varepsilon \in (0, 1/2)$, an ε -significant ensemble at O satisfies $N \leq \exp \left\{ \frac{\left(\frac{A_{\max}(O)}{4\ell_P^2} + h(\varepsilon) \right)}{1 - \varepsilon} \right\}$.

Proof. By Lemma 4.6.2 and minimizing the right-hand side over σ^B , $\ln N \leq \min_{\sigma^B} D_H^\varepsilon(\rho^{XB} \parallel \pi^X \otimes \sigma^B)$. By Lemma 4.6.1 applied with $\rho = \rho^{XB}$, $\sigma = \pi^X \otimes \sigma^B$, this is $\leq \min_{\sigma^B} \left(\frac{[D(\rho^{XB} \parallel \pi^X \otimes \sigma^B) + h(\varepsilon)]}{1 - \varepsilon} \right) = \frac{[\chi + h(\varepsilon)]}{1 - \varepsilon}$ by Lemma 4.6.3, and $\chi \leq S(\bar{\rho}^B)$. Finally, $\bar{\rho}^B$ is itself a state obtained, by (S2'), with no increase in entropy from a bona fide state of the matter crossing the light-sheet bounding $\partial D(O)$; Bousso's theorem (§2.4) applies directly to that light-sheet state, and (S2')'s entropy-non-increase transfers the bound: $S(\bar{\rho}^B) \leq \frac{A_{\max}(O)}{4\ell_P^2}$. Chaining the four inequalities - the last of them now conditional on (S2') - gives the claim. $\angle$

4.6.5 Stage-wise versus historical significance

Theorem 4.6.4 establishes a bound on the cardinality of an ε -significant ensemble that is jointly decodable by a single admissible POVM at a specified measurement stage. Writing $N_\varepsilon(O;t)$ for the largest cardinality of such an ensemble at stage t , the theorem

gives $N_\varepsilon(O;t) \leq \exp \left\{ \frac{\left(\frac{A(t)}{4\ell_P^2} + h(\varepsilon) \right)}{1-\varepsilon} \right\}$. Consequently, whenever $A(t) \leq A_{\max}(O)$, $N_\varepsilon(O;t) \leq \exp \left\{ \frac{\left(\frac{A_{\max}(O)}{4\ell_P^2} + h(\varepsilon) \right)}{1-\varepsilon} \right\}$.

Scope of $A_{\max}$. The proof uses $A_{\max}(O)$ only as an upper bound on the area of the cross-section whose light-sheet carries the matter of (S2'). The theorem therefore also holds with $A_{\max}(O)$ replaced by the stage-local $A_{\max}(O;t)$ of §3.3 (or, where the relevant light-sheet is that of the stage- t cross-section, by $A(t)$ itself). This is the form in which it applies at every finite stage of an observer in an open or flat $\Lambda=0$ FRW cosmology, where the global $A_{\max}(O)$ is infinite.

This is a theorem about jointly distinguishable configurations at a given stage. It should not be identified without further argument with the cardinality of the union of all configurations that may become significant at different stages of O 's history. Indeed, if $C_\varepsilon(O;t)$ denotes the set of configurations jointly ε -decodable at stage t , the historical set $\cup_t C_\varepsilon(O;t)$ may contain configurations associated with different measurement protocols at different stages. A bound on every individual $|C_\varepsilon(O;t)|$ therefore does not, by itself, imply a bound on the historical union.

The passage from the stage-wise operational theorem to the stronger all-of-history formulation is consequently a separate composability problem, stated explicitly in §7(i)(c). This distinction does not weaken Theorem 4.6.4. Rather, it identifies precisely the operational statement that follows from the one-shot argument and prevents an additional historical inference from being hidden inside the definition of $N(O)$.

What this does and does not establish. No postulate of Banks–Fischler type, and no Quantum Focusing Conjecture, appears in the derivation of Theorem 4.6.4. The theorem follows from Bousso's covariant entropy bound together with the stated one-shot quantum-information lemmas, at fixed $\varepsilon > 0$, provided the operational representation assumption of §4.6.0 holds, together with (S2')'s identification of B with an entropy-non-increasing image of the light-sheet-crossing matter - flagged in §4.6.0 and listed as open item (i)(d) below.

Three limitations remain explicit. First, the $\varepsilon \rightarrow 0$ limit is not established by the present hypothesis-testing argument. Second, the classical-quantum representation of a configuration as $X \otimes B$ is used as an operational representation assumption rather than derived as a general theorem of quantum gravity. Third, the stage-wise notion of joint ensemble discrimination does not automatically imply a bound on the union of all configurations that may become significant at different stages over the observer's entire history.

An independent structural check is available from the author's companion note (Materov, 2026a, “The Embedded-Observer No-Go Theorem”). There, Corollary 3.3 gives, for a bipartite closed pure system $H_M \otimes H_{R'}$, the exact identity $I(M:R';t) = 2S(\rho_M(t))$. Reading M as the observer's register $X \otimes B$ and R' as the causal complement of $D(O)$, this identity independently confirms the structural role $S(\rho^B)$ plays above - mutual information (an upper bound on any accessible/Holevo-type information a decoder can extract) is pinned to twice the same entanglement entropy that §2.4/Theorem 4.6.4 bound via $A_{\max}(O)$ - obtained from the author's own entanglement framework rather than from horizon thermodynamics, and landing on the same quantity $S(\rho^B)$ rather than on $\dim H$, in the same spirit as §6's numerical cross-check.

4.7 Dynamical Horizons

§2.4's light-sheet construction is cleanest for a single, static cross-section; Theorem 4.3 as stated requires tracking which cross-section along O 's history actually realizes $A_{\max}(O)$ and confirming the light-sheet construction remains valid throughout a horizon that may be growing or shrinking.

Bousso & Engelhardt (2015a, 2015b) supply exactly this tool. They define a holographic screen as a hypersurface foliated by marginally trapped (or marginally anti-trapped) surfaces - the natural generalization of a single apparent-horizon cross-section to a fully time-dependent setting - and prove, under the null energy condition alone (no staticity, no symmetry assumption), a new area law: the area of successive leaves of a future holographic screen is monotonically non-decreasing (respectively non-increasing along a past holographic screen), independent of any assumption about the sign of $\frac{dA}{dt}$ at any single instant. This is precisely a proof that a well-defined notion of “the maximal cross-section reached so far” exists and behaves monotonically for a fully dynamical, non-static horizon, obtained under the same null energy condition (S2) already assumed for the light-sheet bound itself.

Combining this with §3.3's apparent-horizon reading of $A_{\max}(O)$: under (S2), $A_{\max}(O)$ is realized as the supremum of a monotonically evolving quantity along a holographic screen adapted to O 's worldline, rather than as an ad hoc supremum over an unstructured set of cross-sections, and Bousso's light-sheet bound applies leaf-by-leaf along the screen by construction. This does not remove every technical subtlety - in particular, a transition between a future and a past holographic screen (as can occur, e.g., across a bounce or a horizon-crossing event) still requires separate bookkeeping, which is not worked out here - but it replaces an unaddressed difficulty with a cited, proved monotonicity result covering the generic case.

5. Complementarity with the Causality Principle

The causality principle and the Principle of Significance are structurally complementary, not merely thematically parallel. Causality is a statement about the causal structure of spacetime alone - it constrains propagation and holds regardless of any matter content or observer's resources. Significance, as formalized here, is a statement about entropy and information — it constrains capacity, and is sensitive to the matter content (through the null energy condition) and to the observer's resources (through the operational reading of Definition 3.2). One restricts the geometry of what can reach a point; the other restricts what a bounded recipient at that point can ever make of what reaches it.

6. A Numerical Cross-Check: Two Independent Routes to 10^{120}

The de Sitter horizon entropy of our own universe, computed via the Gibbons–Hawking area law from the observed value of the cosmological constant, is of order 10^{120} – 10^{122} in Planck units. Lloyd (2002) arrives at a strikingly similar figure - no more than 10^{120} elementary operations on 10^{90} bits (10^{120} bits including gravitational degrees of freedom) over the universe's history - by a wholly unrelated argument: the Margolus–Levitin bound combined with an entropy bound on registerable bits, applied to matter and energy content rather than horizon geometry. The two routes' landing in the same order of magnitude is a nontrivial, if informal, consistency check, and is not evidence for any specific value of $N(O)$ in Theorem 4.3, which remains conditional on (S3b).

7. Open Problems

(i) Four items Theorem 4.6.4 and Definition 3.2' leave open. §4.6 derives, rather than conjectures, a bound of Bousso's form on $N(O)$, but at fixed error tolerance $\epsilon > 0$, conditional on the $X \otimes B$ representation of “configuration,” and conditional on Definition 3.2' being the right formalization of significance accumulated over an observer's history. Four precise technical items remain:

(a) whether the $\epsilon \rightarrow 0$ limit can be taken rigorously to recover an exact, error-free statement - the $\epsilon = 0$ hypothesis-testing entropy is governed by state supports rather than by $D(\rho|\sigma)$, so Lemma 4.6.1's proof does not directly extend;

(b) whether the classical-register representation of a “configuration” is physically forced for the degrees of freedom crossing a light-sheet, or is an additional structural assumption;

(c) Historical composability of operational significance. Theorem 4.6.4 bounds the cardinality of an ϵ -significant ensemble that is jointly decodable by a single POVM at a specified measurement stage. The stronger historical quantity $|\cup_t C_\epsilon(O;t)|$ is not automatically bounded by the supremum of the individual $N_\epsilon(O;t)$. The reason is that the POVM realizing significance may depend on t . Configurations that are distinguishable at different stages need not therefore form a single jointly decodable ensemble.

A complete derivation of the original all-of-history Principle of Significance would require an additional composability or information-accumulation theorem controlling the union of operationally distinguishable configuration sets over the observer's worldline. There are at least two possible routes. One would show that the historical union can itself be embedded into a single bounded-capacity discrimination problem. Another would replace the historical cardinality N_{hist} by an additive operational information measure whose total accumulation is directly bounded by the relevant geometric entropy budget. Until such a result is established, the all-of-history reading of Theorem 4.3 remains conditional. The fixed-stage operational bound of Theorem 4.6.4, however, is unaffected by this limitation.

(d) Light-sheet realization of B. The proof of Theorem 4.6.4 applies the covariant entropy bound to $\bar{\rho}_B$, the averaged state on the operationally-defined system B of §4.6.0, which characterizes B only as the degrees of freedom the observer's measurement acts on. Nothing there establishes that B is, or is obtained without increasing entropy from, the matter crossing the light-sheet bounding $\partial D(O)$, which is what Bousso's theorem directly bounds. Assumption (S2') makes this identification explicit as a further working hypothesis; establishing it - or the weaker entropy-non-increase $S(\bar{\rho}_B) \leq S(\rho_{LS})$ alone - for a general physical realization of B , rather than positing it, is left open here. Note that the natural candidate route - simply inclusion, $B \subseteq LS$ with the rest traced out - does not work in general: for LS in a pure, highly entangled state, tracing out a subsystem can only increase entropy above

the (vanishing) entropy of the whole, so $S(\rho_B) \leq S(\rho_{LS})$ is not a generic property of subsystems and is not claimed to be one here. What is needed instead is physically natural sufficient conditions on the specific coupling/measurement channel realizing B - e.g. restrictions on the entanglement of the relevant light-sheet matter with the rest of the light-sheet, or on the class of measurement processes admitted - under which entropy non-increase can be established rather than assumed.

All four are now sharply stated, self-contained problems in one-shot quantum information theory and quantum gravity, rather than a single diffuse conjecture.

(ii) Finiteness of $A_{\max}$ beyond the cases of §3.3. §3.3 gives global finiteness only where the apparent-horizon area stays bounded along O's history, and handles the $\Lambda=0$ open or flat FRW case with unbounded history at the stage-local level only (there the global maximal area is infinite and Theorem 4.3 does not apply); the same construction is not automatic for spacetimes lacking FRW symmetry or a well-defined global time slicing (e.g. strongly inhomogeneous or genuinely rotating spacetimes), where the marginally-trapped-surface construction underlying the apparent/holographic-screen horizon can be foliation-dependent or fail to exist globally. Establishing (S1) in that generality remains open.

(iii) Screen transitions. §4.7's monotone area law is proved separately for future and past holographic screens; a fully general statement covering transitions between the two (e.g. at a cosmological bounce) is, to the author's knowledge, not yet available in the literature cited and is not attempted here.

(iv) The algebraic Holevo bound. Whether an intrinsic, algebra-only N-ary discrimination bound holds for a general observer's own Type II von Neumann algebra of observables A_O - tied to $A_{\max}(O)/4\ell_P^2$ with no external $X \otimes B$ register - replacing the operational representation assumption of §4.6.0 entirely. Resolved: §9.6–9.7 (Appendix) derive this bound from a published theorem (Podsiadkowska & Wieczorek, 2017) plus classical Fano, with the additive constant relating their Segal entropy to the CLPW/JSS generalized entropy identified and verified against CLPW's primary text. What remains open is not this information-theoretic chain but Jensen–Sorce–Speranza's own separate Assumption A4 on the general existence of a suitable A_O for an arbitrary bounded observer (§9.2 of the Appendix) - now confirmed for Killing horizons (Faulkner & Speranza, 2024; Ali & Suneeta, 2024), open in the fully general case.

7.1 Worked Example (Flash and blast wave). Why naive stage-summation overcounts historical significance - a concrete two-stage scenario.

Let O observe two causally ordered informative events along its worldline, both originating from a single distant energetic event: at stage t_1 , O detects electromagnetic radiation (“the flash”) from the event; at a later stage $t_2 > t_1$, O is causally reached by a slower-propagating disturbance (“the blast wave”) originating from the same event. Since the flash propagates at c and the leading edge of the blast wave propagates at some $v < c$, the interval $t_2 - t_1 > 0$ is set by the difference in propagation time over the relevant distance.

Let X_1 be a random variable ranging over the N_1 “flash configurations” (source type, magnitude, direction, etc.) distinguishable at stage t_1 in the sense of Definition 3.2', so $|C_e(O; t_1)| \leq N_1$ by Theorem 4.6.4. Let X_2 range over the N_2 “blast configurations” (arrival profile, residual energy distribution, etc.) similarly distinguishable at t_2 : $|C_e(O; t_2)| \leq N_2$.

The naive composition, and why it overcounts

One might guess that the number of jointly distinguishable two-stage histories available to O is bounded by the product $N_1 N_2$ - as if X_1 and X_2 were independently drawn, one full “budget” of distinguishability per stage. This is exactly the kind of composition §4.6.5 warns is *not* licensed by Theorem 4.6.4, and this example makes the reason concrete: X_2 is not independent of X_1 . Both are generated by the same source event, so a substantial part of what X_2 will turn out to be is already implicit in X_1 - the blast profile is, to a good approximation, a function of the same source parameters the flash already revealed.

Formally: the number of jointly distinguishable two-stage histories is bounded not by $N_1 \cdot N_2$ but by

$$\exp[H(X_1) + H(X_2|X_1)] = \exp[H(X_1, X_2)] = \exp[H(X_1) + H(X_2) - I(X_1; X_2)],$$

which is $\leq N_1 \cdot N_2$ with equality only when $I(X_1; X_2)=0$ (the two stages genuinely independent). In this scenario $I(X_1; X_2)$ can be large - in the limiting case where the blast profile is a deterministic function of the flash data via known physics, $H(X_2|X_1)=0$ and the joint bound collapses to N_1 alone: stage t_2 contributes essentially *no* new distinguishable content beyond what t_1 already fixed, even though $|C_e(O; t_2)|$ itself - the stage-local bound Theorem 4.6.4 gives at t_2 in isolation - may still be large.

What this does and does not show

This does *not* show that the interval between t_1 and t_2 carries zero significance for O. O remains within an actively unfolding causal patch, and the specific manner of the blast's arrival - timing fluctuations, details not fixed by the coarse flash data, the response of whatever systems remain - still carries some residual amount of operationally distinguishable content: $H(X_2|X_1)$ in

general, though reduced, need not vanish. The precise claim is narrower and entirely quantitative: *marginal, newly accruing* significance drops sharply once a common cause has already been resolved, and any composability theorem for $N(O)$ across O 's history must track this via conditional entropy or mutual information - not by summing or multiplying stage-local bounds as though successive stages were independent.

This connects to, without depending on, the companion No-Go theorem's Corollary 3.3 identity $I(M:R';t) = 2S(\rho_M(t))$: both results track the same qualitative fact - that information already acquired through correlation with a common source reduces what remains genuinely new to learn - in two different registers, the quantum-entanglement register there and the classical operational register (Definition 3.2') here.

A terminological note. “Significance” in this example is used exclusively in Definition 3.2's operational sense -distinguishability by a measurement protocol - and carries no evaluative or existential connotation. An observer's assessment that some later stage of its history matters less to it in some further, evaluative sense is a separate question this framework neither models nor needs to. That is a task for philosophy.

8. Relation to Prior Work

This note's central technical contribution is the derivation in §4.6: a Bousso-form bound on $N(O)$ obtained from Bousso's covariant entropy bound and elementary one-shot quantum information theory alone, with no postulate of Banks–Fischler or Quantum-Focusing-Conjecture type anywhere in the chain - at the cost of a fixed error tolerance $\epsilon > 0$, an explicit representation assumption, and the light-sheet realization assumption (S2'), all stated rather than hidden. §4.5 separately clarifies the relationship between this operational bound and the stronger, purely ontological Banks–Fischler proposal, which the derivation does not require. §3.3 and §4.7 address the finiteness-of- $A_{\max}$ and dynamical-horizon difficulties using apparent-horizon thermodynamics and the Bousso-Engelhardt area law respectively (global finiteness holds only where the apparent-horizon area stays bounded; otherwise only a stage-local version is finite), together with one internal cross-check drawn from the author's companion “The Embedded-Observer No-Go Theorem” paper. The Appendix (§9) extends the derivation to an intrinsic, algebra-only formulation, replacing the $X \otimes B$ representation assumption with the observer's own algebra of observables, following Chandrasekaran-Longo-Penington-Witten's and Jensen-Sorce-Speranza's crossed-product constructions.

9. Toward an Algebraic Formulation (Appendix)

A scoped program for replacing the operational representation assumption of §4.6.0 with the observer's intrinsic algebra of observables.

9.1 Motivation: Replacing $X \otimes B$ with A_O

§4.6.0's operational representation assumption - that the discrimination task admits a classical-quantum representation on $X \otimes H_B$ for some system B carrying “precisely the degrees of freedom on which the observer's admissible measurement acts” - is stated honestly as an assumption in §4.6, but it purchases the derivation of Theorem 4.6.4 at the cost of positing a tensor-factor structure that quantum gravity does not obviously supply on its own. A more intrinsic route would define the observer's information capacity directly from its own algebra of observables A_O , without positing any such factorization from outside. This appendix states precisely how far that route can currently be walked, and where it stops.

9.2 Background: From Type III to Type II

Ordinary quantum field theory assigns to a causal diamond, wedge, or static patch a von Neumann algebra of local observables that is generically a Type III₁ factor. Type III₁ algebras admit no trace, hence no density matrices and no well-defined entropy for individual states - the mathematical root of the divergences that plague naive attempts to assign an entanglement entropy to a QFT subregion.

Witten (2022) showed that including gravity changes this: enlarging the algebra by a *crossed product* with the modular automorphism group, generated physically by a bounded-below Hamiltonian (the observer's own clock/energy degree of freedom), converts the Type III₁ algebra into a Type II algebra, which does admit a trace and hence a well-defined, renormalized entropy for each state.

Chandrasekaran, Longo, Penington & Witten (2022/2023, CLPW) carried this out explicitly for a geodesic observer in the static patch of de Sitter space: the resulting algebra A_O is Type II₁, possesses a maximum-entropy state (corresponding to empty de Sitter space), and the entropy of any semiclassical state agrees, up to a state-independent additive constant, with the generalized entropy $S_{\text{gen}} = \frac{A}{4G} + S_{\text{out}}$. §9.7 below determines this constant explicitly.

Jensen, Sorce & Speranza (2023, JSS) generalized this construction from the de Sitter static patch to general spatially compact subregions of a spacetime satisfying Einstein's equations coupled to matter (in the $G_N \rightarrow 0$ limit), introducing an explicit model of an observer living inside the region. They again obtain a Type II₁ algebra - explicitly linking the existence of its entropy-maximizing state to “a version of Jacobson's entanglement equilibrium hypothesis” - with entropy agreeing with the generalized entropy up to a constant. This is precisely the generalization needed to cover an arbitrary bounded observer O in the sense of Definition 3.1, rather than only the highly symmetric de Sitter case. They are explicit, however, that the construction rests on their own Assumption A4 (§6.3 of their paper): the existence of a state whose modular flow is geometric at an instant in time, near a spacelike Cauchy surface - an assumption not established in general, and logically independent of the additive-constant question resolved in §9.7.

Status of Assumption A4. Faulkner & Speranza (2024) prove that such states are guaranteed to exist for Killing horizons, in both free and interacting quantum field theories, via the properties of half-sided translations on the horizon (the same mechanism underlying Borchers' theorem for half-sided modular inclusions) - giving a genuine, non-conjectural construction of A_O for this class of horizons. Ali & Suneeta (2024) develop a related, local-horizon version of the construction that includes explicit observer degrees of freedom. Both results confirm A4 for Killing horizons specifically; a fully general proof, for a spacelike Cauchy surface with no Killing symmetry, remains open. A closely related but distinct question - uniqueness of *null-local* (rather than spacelike-Cauchy-surface) modular flow - is examined by Satishchandran & Sorce (2026), who find a more restrictive result (the vacuum sector admits only a trivial such flow); this bears on, but does not settle, A4 as JSS state it.

9.3 The Target Statement

Write A_O for the observer's algebra as constructed, where the construction is currently established, by CLPW or JSS; $S(A_O)$ for the entropy of a state on this algebra in their sense; and $D_{\text{Araki}}(\cdot|\cdot)$ for the Araki (1976) relative entropy of two states on a von Neumann algebra, defined via the relative modular operator of Tomita -Takesaki theory. Araki's relative entropy is known to satisfy monotonicity under restriction to subalgebras - precisely the data-processing inequality on which the proof of Lemma 4.6.1 depends.

Conjecture 9.3.1 (Algebraic Principle of Significance). Let O be a bounded observer whose algebra A_O is Type II in the sense of CLPW/JSS. For any $\varepsilon \in (0, \frac{1}{2})$ and any ensemble of N states $\{\omega_1, \dots, \omega_N\}$ on A_O jointly ε -discriminable by a single POVM affiliated to A_O (in the sense of an algebraic analogue of Definition 3.2'),

$$\ln N \leq \frac{[S(A_O) + h(\varepsilon)]}{1 - \varepsilon},$$

with $S(A_O) \lesssim \frac{A_{\text{max}}(O)}{4\ell_P^2}$ by the CLPW/JSS entropy-area relation. §9.7 establishes this inequality with $S(A_O)$ identified precisely.

9.4 What Would Be Needed to Prove This

Establishing Conjecture 9.3.1 by the same route as Theorem 4.6.4 requires three ingredients (numbered to avoid collision with §7's open-problem lettering):

Ingredient 1 [Established]. A well-defined relative entropy for states on a general von Neumann algebra, satisfying the data-processing inequality - supplied by Araki (1976) and, in the specific form needed below, Uhlmann (1977).

Ingredient 2 [Established, conditionally]. A_O is Type II with a canonical entropy-maximizing state - supplied by CLPW (2022/2023) for a de Sitter geodesic observer, and by JSS (2023) for a general spatially compact subregion, conditional on JSS's own Assumption A4 (§9.2): existence of a state with instantaneously geometric modular flow near a spacelike Cauchy surface. Confirmed for Killing horizons (Faulkner & Speranza, 2024; Ali & Suneeta, 2024); open in general. The maximal-state / relative-entropy structure that follows once A4 holds is independently corroborated as a general feature of this construction by Witten (2024) and Blommaert, Kudler-Flam & Urbach (2025).

Ingredient 3 [Resolved]. An algebraic, intrinsic N-ary discrimination bound for states and POVMs on a Type II von Neumann algebra, with no external $X \otimes B$ register, and with its final step - relating the bound to $\frac{A_{\text{max}}(O)}{4\ell_P^2}$ - resting on a correctly identified quantity rather than a hopeful sign guess. The author raised the general question directly with E. Witten, one of CLPW's authors, who replied that he did not have a useful suggestion - confirming this was a genuine gap unaddressed by the gravity/algebraic-QFT community that produced Ingredient 2, not an oversight of the present literature search. §9.6 supplies the discrimination bound itself (Proposition 9.6.1), via Podsiadkowska & Wieczorek's (2017) published Generalized Holevo Theorem combined with the classical Fano inequality; §9.7 supplies the final step, after an initial sign error was caught, corrected, and then verified against CLPW's primary text directly.

9.5 A Proof Sketch for Ingredient 3, Part One: The Hypothesis-Testing Step (Secondary Route)

This subsection reports the author's own, independent attempt at closing the Lemma 4.6.1 part of Ingredient 3, by a different route from §9.6/§9.7. It is presented as a proof sketch, not a verified theorem - it has not been checked by a specialist in operator algebras, and the author is not one. It is retained, despite §9.6/§9.7 now closing this ingredient by a more solid route, because it isolates a self-contained, potentially independently useful lemma and because exposing it in full lets any error be checked line by line.

Definition 9.5.1 (Algebraic hypothesis-testing relative entropy). Let M be a von Neumann algebra and ω_ψ, ω_η normal states on M . For $\varepsilon \in (0,1)$, define

$$D_H^\varepsilon(\omega_\psi \| \omega_\eta) := -\ln \inf \{ \omega_\eta(Q) : Q \in M, 0 \leq Q \leq 1, \omega_\psi(Q) \geq 1-\varepsilon \}.$$

Unlike the finite-dimensional definition used in §4.6, this makes no reference to a trace, a Hilbert-space tensor factorization, or a density matrix - it is meaningful for any von Neumann algebra, Type II included, since it is phrased entirely in terms of positive elements of M and states on M .

Lemma 9.5.2 (Algebraic one-shot converse - author's proof sketch, not independently verified). Let M be a von Neumann algebra in standard form, ω_ψ, ω_η normal states with ω_η faithful, and $\varepsilon \in (0, 1/2)$. Then

$$D_H^\varepsilon(\omega_\psi \| \omega_\eta) \leq \frac{[S(\omega_\psi | \omega_\eta) + h(\varepsilon)]}{1-\varepsilon},$$

where $S(\cdot | \cdot)$ is the Araki (1976) relative entropy.

Proof sketch.

Step 1 (existence of a minimizer). The set $\{Q \in M : 0 \leq Q \leq 1\}$ is weak-* compact by the Banach-Alaoglu theorem, since M is a dual Banach space (the predual of a von Neumann algebra always exists). The constraint $\omega_\psi(Q) \geq 1-\varepsilon$ is weak-* closed since ω_ψ is normal, hence weak-* continuous; the intersection of a weak-* compact set with a weak-* closed set is weak-* compact. The functional ω_η is likewise weak-* continuous, so it attains its infimum on this compact set at some Q^* .

Step 2 (reduction to a classical binary measurement). Q^* defines a normal, unital, completely positive map $T: M \rightarrow \mathbb{C}^2$, $T(f) = f(1)Q^* + f(2)(1-Q^*)$ for $f \in \mathbb{C}^2$ - i.e. the two-outcome POVM $\{Q^*, 1-Q^*\}$. Under the dual (Schrödinger-picture) action, T sends ω_ψ, ω_η to classical distributions $p=(p_1, p_2)=(\omega_\psi(Q^*), 1-\omega_\psi(Q^*))$ and $q=(q_1, q_2)=(\omega_\eta(Q^*), 1-\omega_\eta(Q^*))$.

Step 3 (data-processing inequality). By the DPI for the Araki relative entropy on general von Neumann algebras (Araki, 1976; Uhlmann, 1977; see also Hollands, 2021, eq. (1)), applied to T : $S(\omega_\psi | \omega_\eta) \geq S(\omega_\psi \circ T | \omega_\eta \circ T) = S(p | q)$, where $S(p | q)$ is the ordinary classical Kullback-Leibler divergence.

Step 4 (classical algebra, verbatim as in Lemma 4.6.1). By construction $p_1 = \omega_\psi(Q^*) \geq 1-\varepsilon$. Exactly as in the proof of Lemma

4.6.1: $S(p | q) \geq -H(p) - p_1 \ln q_1$, so $-\ln q_1 \leq \frac{[S(\omega_\psi | \omega_\eta) + h(\varepsilon)]}{1-\varepsilon}$, and since $q_1 = \omega_\eta(Q^*)$ is exactly the defining minimum, $D_H^\varepsilon(\omega_\psi \| \omega_\eta) = -\ln q_1$, giving the claim. $\square$ (sketch)

Status. Ingredient 1 and the classical mechanism used here transfer cleanly to any von Neumann algebra, as this sketch shows. What it does not supply, and what §9.6 supplies instead by a different and more solid route, is the N-ary ensemble bound and its connection to $A_{\max}(O)$. This sketch remains unverified by a specialist and is offered for scrutiny in its own right, independent of whether §9.6–§9.7 are accepted.

9.6 A Cleaner Route: Fano's Inequality and the Podśędkowska–Wieczorek Theorem

Podśędkowska & Wieczorek (2017) supply the N-ary bound directly, by an argument that does not go through Lemma 9.5.2's hypothesis-testing route at all, resting entirely on their published, peer-reviewed Theorem 1 rather than on an unverified sketch.

Their setting matches Definition 3.2' directly. Podśędkowska & Wieczorek work with states $\rho_1, \dots, \rho_N$ on a von Neumann algebra M with a finite faithful normal trace τ , occurring with prior probabilities π_i , and a measurement $M=(M_1, \dots, M_m)$ with $M_j \in M$, $M_j > 0$, $\sum M_j = 1$ - a POVM affiliated to M directly, with no external classical register X or auxiliary system B posited. This is exactly the intrinsic, algebra-only formulation §9.1 wanted in place of §4.6.0's $X \otimes B$ representation assumption.

Theorem (Podśędkowska & Wieczorek, 2017, Theorem 1). $I_{\text{acc}} \leq \chi$, where $I_{\text{acc}} := \max_M I(\rho, M)$ is the accessible information (mutual information of the joint distribution $p(i, j) = \pi_i \rho_i(M_j)$, maximized over measurements), and $\chi := S(\sum \pi_i \rho_i) - \sum \pi_i S(\rho_i)$ is the Holevo χ -quantity, with S the Segal entropy on (M, τ) . Their proof (§3.1) chains strong subadditivity of Segal's entropy on a

general finite von Neumann algebra (their Lemma 1, extending Lieb's convex trace-function technique via Epstein's generalization to C^* -algebras) with the data-processing property of relative entropy under normal Schwarz maps (their Lemma 2, citing Neshveyev–Størmer). Both are established, published results.

9.6.1 Proposition (Algebraic N-ary bound via Fano). Let $M=A_O$ be a von Neumann algebra with finite faithful normal trace τ . Let $\{\rho_1, \dots, \rho_N\}$ be states on M , uniformly weighted, jointly ε -significant in the sense of Definition 3.2' (a POVM $\{M_i\}$ affiliated to M , $\sum M_i \leq 1$, with $\frac{1}{N} \sum \rho_i(M_i) \geq 1 - \varepsilon$). Then

$$\ln N \leq \frac{[\chi + h(\varepsilon)]}{1 - \varepsilon},$$

Proof. Let X be uniform on $\{1, \dots, N\}$ and Y the outcome of the decoder POVM $\{M_i\}$. By ε -significance, the average success probability $\frac{1}{N} \sum \rho_i(M_i) \geq 1 - \varepsilon$. By the classical Fano inequality, $H(X|Y) \leq h(\varepsilon) + \varepsilon \ln(N-1) \leq h(\varepsilon) + \varepsilon \ln N$, so $I(X; Y) \geq (1 - \varepsilon) \ln N - h(\varepsilon)$. Since $I(X; Y) = I(\rho, M) \leq I_{\text{acc}} \leq \chi$ (by the theorem above), chaining gives $\ln N \leq \frac{[\chi + h(\varepsilon)]}{1 - \varepsilon}$. ■

Nothing in this proof is original beyond the assembly: the Fano step is elementary and standard (Cover & Thomas, 2006, Thm. 2.10.1), and the theorem it is chained to is Podsiędkowska & Wieczorek's own published result.

9.7 Closing the Loop: Identifying the Additive Constant

Proposition 9.6.1 bounds N in terms of χ , not yet in terms of $\frac{A_{\text{max}}(O)}{4\ell_p^2}$. Closing that gap requires knowing what χ is bounded by. An initial attempt at this final step contained a sign error, caught and corrected through direct verification against Podsiędkowska & Wieczorek's own text and then against CLPW's primary source; the corrected argument is given here in full, since the error is instructive and the correction is not obvious without it.

9.7.1 The obstruction. Segal's entropy, as defined by Podsiędkowska & Wieczorek (following Segal, 1960), satisfies $S(\rho) \leq \tau(1) - \tau(\hat{\rho})$ - an *upper* bound, not a lower one. For $\tau(1)=1$ (the Type II_1 case), this gives $S(\rho) \leq 0$ for every state, with equality only at the tracial state $\rho=\tau$. This is elementary: writing $H(\rho) := \tau(\hat{\rho} \log \hat{\rho})$, the inequality $t \log t \geq t-1$ (true for all $t>0$, minimum at $t=1$) gives $H(\rho) = \int t \log t \, d\tau(e(dt)) \geq \int (t-1) d\tau(e(dt)) = \tau(\hat{\rho}) - \tau(1) = 0$, so $S(\rho) = -H(\rho) \leq 0$. Consequently $\chi = S(\bar{\rho}) - \sum \pi_i S(\rho_i) \geq S(\bar{\rho})$ trivially (subtracting non-positive terms), which is the *wrong* direction for bounding χ from above, and $S(\bar{\rho}) \leq 0$ itself gives no useful finite ceiling.

9.7.2 The resolution. Since τ is tracial, its own density is the identity, so for any state ρ : $D(\rho|\tau) = \tau(\hat{\rho} \log \hat{\rho}) - \tau(\hat{\rho} \log 1) = H(\rho) = -S(\rho)$. That is, Segal's entropy is exactly minus the Araki relative entropy to the tracial state: $S(\rho) = -D(\rho|\tau)$. Using this identity, $\chi = S(\bar{\rho}) - \sum \pi_i S(\rho_i) = \sum \pi_i D(\rho_i|\tau) - D(\bar{\rho}|\tau) \leq \sum \pi_i D(\rho_i|\tau)$, since $D(\bar{\rho}|\tau) \geq 0$ (Araki/Uhlmann positivity). This is now an upper bound in the correct direction, expressed via a manifestly non-negative quantity.

9.7.3 Confirmation from CLPW's primary text. Reading Chandrasekaran, Longo, Penington & Witten (2023) directly confirms this is not a reinterpretation but their own stated convention: “ $S(\rho_A)$ is the same as the entropy deficit $\Delta S_A = S(\rho_A) - S(\rho)$, since $S(\rho)=0$,” where $\rho=1$ is their maximum-entropy (empty de Sitter) state - i.e. their S is already, by construction, an entropy *deficit* relative to the maximal state, matching $S(\rho) = -D(\rho|\tau)$ exactly. Their stated relation $S_{\text{CLPW}}(\psi) = S_{\text{gen}}(\psi) + \text{const}$, combined with $S_{\text{CLPW}} \leq 0$ with maximum 0 at $\rho=\tau$, pins the additive constant precisely: $\text{const} = -A_{\text{ds}}/4G_N$, the *maximum* generalized entropy (that of empty de Sitter itself, the entropy-maximizing state). Hence

$$D(\rho|\tau) = -S_{\text{CLPW}}(\rho) = \frac{A_{\text{ds}}}{4G_N} - S_{\text{gen}}(\rho).$$

Since $S_{\text{gen}}(\rho) \geq 0$ always - an elementary, essentially definitional fact about generalized entropy, needing none of the covariant-bound machinery of §2.4 - this gives $D(\rho|\tau) \leq \frac{A_{\text{ds}}}{4G_N}$ directly.

9.7.4 Generalization beyond de Sitter. The same structure is not a peculiarity of CLPW's specific symmetry. It follows from the general Murray–von Neumann classification: *any* Type II_1 factor has a canonical entropy-maximizing (tracial) state, so the identity $S = -D(\cdot|\tau)$ and the resulting bound hold for any Type II_1 construction, not only CLPW's. This is confirmed directly by JSS's own abstract, which states that for spatially compact regions their algebra is Type II_1 , “implying the existence of an entropy maximizing state, which realizes a version of Jacobson's entanglement equilibrium hypothesis” — the same structure, now for a general bounded subregion rather than only the de Sitter static patch. It is corroborated further by two later papers that make the identification of entropy with relative entropy to a maximal state fully explicit and general: Witten (2024, “A background-independent algebra in quantum gravity”) writes that “it is natural to think of the Hartle–Hawking no-boundary state as a universal state of maximum entropy, and to define entropy in terms of the relative entropy with this state”; and Blommaert, Kudler-Flam & Urbach (2025) show the observer's no-boundary state is the identity operator (the maximum-entropy state) and

“allow one to define von Neumann entropy for observers in different universes as the relative entropy with respect to the no-boundary state.” Between CLPW's original computation and these two independent confirmations, the pattern $D(\rho|\tau) \leq$ (maximum generalized entropy available to O) is now a recurring, not an idiosyncratic, feature of this line of work.

This is logically independent of JSS's own separate open conjecture (the existence of a state with instantaneously geometric modular flow, §9.2): that conjecture concerns *whether* A_O is Type II₁ with the stated entropy formula at all for a general bounded observer; §9.7's result concerns what follows *once* it is.

9.7.5 The final chain. Combining §9.6 and §9.7: writing $A_{\max}(O)$ for the maximum generalized entropy (in area units) available to O's algebra - A_{ds} in the CLPW case, the analogous maximal-state value in the JSS case:

$$\chi \leq \Sigma \pi_i D(\rho_i | \tau) \leq \frac{A_{\max}(O)}{4\ell_P^2},$$

and hence, via Proposition 9.6.1,

$$\ln N \leq \frac{\left[\frac{A_{\max}(O)}{4\ell_P^2} + h(\varepsilon) \right]}{1-\varepsilon},$$

which is exactly Conjecture 9.3.1's target inequality. Ingredient 3 is accordingly resolved, modulo ordinary due diligence (an independent check of §9.7.1-9.7.3's algebra by a specialist would still be worthwhile) and modulo JSS's own separate, already-identified conjecture about the general existence of A_O itself (§9.2, Ingredient 2).

9.8 Status

This appendix began as a precisely stated program with all three ingredients open or merely candidate-supported. It now stands considerably more resolved: Ingredient 1 was always established; Ingredient 2 is established for CLPW and conditionally for JSS, with the conditionality now understood to concern only the existence and identification of A_O itself, not the entropy-bound mechanics; Ingredient 3 is resolved by §9.6-§9.7, via a route that survived a real sign error, an honest rollback, and a primary-source re-verification - rather than by the author's own §9.5 sketch, which is retained only as an independent, unverified secondary route. Item (iv) in §7 reflects this.

A note on how this was found. The author put the general question of Ingredient 3 directly to E. Witten, one of CLPW's authors, whose response was that he did not have a useful suggestion. This reflects that the gravitational-algebra community's expertise sits in QFT and operator-algebra methods, not one-shot quantum information theory specifically - the pieces that closed Ingredient 3 (Podsiadkowska & Wieczorek, 2017; the Fano inequality) come from exactly the adjacent field that community would have no particular reason to have tracked. The final obstruction (§9.7.1) was not a literature gap at all, but an error in the author's own reasoning about a quantity CLPW define explicitly in their own paper - caught only by reading that paper's stated results line by line rather than reasoning from its abstract and secondary characterizations. Both morals are worth stating plainly for anyone extending this work further: check adjacent subfields before assuming a gap is real, and check primary sources before assuming a fix works.

Status Summary

Component	Status	Depends on A4?	Source
Covariant entropy bound on light-sheets (Prop. 4.1)	Established under stated hypotheses (conjectured in Bousso 1999; see §4)	No	Bousso (1999); Flanagan, Marolf & Wald (2000); Bousso et al. (2014)
Cosmological / apparent horizon area law	Theorem (semiclassical gravity)	No	Gibbons & Hawking (1977); Cai & Kim (2005)
(S3a) $\dim H = \exp(S)$ -ontological	Conjecture - contested, not needed	No	Banks (2000); Fischler (2000)
(S3b) operational distinguishability $\leq \exp(S)$	Derived at fixed $\varepsilon > 0$ from Bousso's theorem + one-shot QI, conditional on (S2') (Thm. 4.6.4); $\varepsilon \rightarrow 0$ limit, the $X \otimes B$ representation, the (S2') light-sheet realization, and the ensemble-vs-history equivalence remain open - see §7(i)	No	§4.6; Wang & Renner (2012); Nielsen & Chuang (2000)
Finiteness of $A_{\max}$ for $\Lambda=0$ FRW	Stage-local: finite (§3.3). Global: only where $A_{\max}$ the apparent-horizon area is bounded; open for unbounded open/flat histories (§7(ii))	No	§3.3; Cai & Kim (2005)
Dynamical horizons	Resolved via a proved monotone area law	No	§4.7; Bousso & Engelhardt (2015a, 2015b)

Principle of Significance (Thm. 4.3)	Conditional theorem; historical (all-of-history) reading further conditional on §7(i)(c)	Only via row below (A(O) Type II), for its algebraic reading - the stage-wise (§4.6) reading does not use A4 at all	This paper
A(O) is Type II with entropy-maximizing state	Established for de Sitter; general case conditional on JSS's Assumption A4; confirmed for Killing horizons; the maximal-state structure itself confirmed general	Yes -this row is Assumption A4 itself	CLPW (2022/2023); Jensen, Sorce & Speranza (2023); Faulkner & Speranza (2024); Ali & Suneeta (2024); Witten (2024); Blommaert, Kudler-Flam & Urbach (2025)
Algebraic N-ary bound (Prop. 9.6.1)	Derived from a published theorem + classical Fano - holds for any Type II von Neumann algebra with finite faithful trace	No -self-contained; meaningful for A_O only once the row above holds	§9.6; Podędkowska & Wieczorek (2017)
Additive constant / final step (§9.7)	Resolved -identified as the maximum generalized entropy available to O, via $S = -D(\cdot \tau)$; holds for any Type II _i factor	No -self-contained; meaningful for A_O only once the Type II row above holds	§9.7; verified against CLPW primary text
Lemma 4.6.1 analogue (secondary route)	Author's proof sketch (Lemma 9.5.2) - unverified, not needed given §9.6–§9.7	No	§9.5; Araki (1976); Uhlmann (1977)

References

- Ali, M., & Suneeta, V. (2024). A local generalized second law in crossed product constructions. arXiv:2404.00718.
- Araki, H. (1976). Relative entropy of states of von Neumann algebras. *Publications of the Research Institute for Mathematical Sciences (Kyoto)*, 11(3), 809–833.
- Bak, D., & Rey, S.-J. (2000). Cosmic holography. *Classical and Quantum Gravity*, 17(15), L83.
- Balakrishnan, S., Faulkner, T., Khandker, Z. U., & Wang, H. (2019). A general proof of the quantum null energy condition. *Journal of High Energy Physics*, 2019(9), 020 (arXiv:1706.09432).
- Banks, T. (2000). Cosmological breaking of supersymmetry, or little Lambda goes back to the future II. *Characterized in Bousso* (2002).
- Bao, N., & Ooguri, H. (2017). Distinguishability of black hole microstates. *Physical Review D*, 96, 066017 (arXiv:1705.07943).
- Bekenstein, J. D. (1973). Black holes and entropy. *Physical Review D*, 7, 2333–2346.
- Bekenstein, J. D. (1981). A universal upper bound on the entropy-to-energy ratio for bounded systems. *Physical Review D*, 23, 287–298.
- Blommaert, A., Kudler-Flam, J., & Urbach, E. Y. (2025). Absolute entropy and the observer's no-boundary state. *Journal of High Energy Physics*, 2025(11), 113 (arXiv:2505.14771).
- Bousso, R. (1999). A covariant entropy conjecture. *Journal of High Energy Physics*, 9905, 004 (arXiv:hep-th/9905177).
- Bousso, R. (2002). Adventures in de Sitter space (arXiv:hep-th/0205177).
- Bousso, R., Casini, H., Fisher, Z., & Maldacena, J. (2014). Proof of a quantum Bousso bound. *Physical Review D*, 90, 044002 (arXiv:1404.5635).
- Bousso, R., & Engelhardt, N. (2015a). A new area law in general relativity. arXiv:1504.07627.
- Bousso, R., & Engelhardt, N. (2015b). Proof of a new area law in general relativity. *Physical Review D*, 92, 044031 (arXiv:1504.07660).
- Bousso, R., Fisher, Z., Koeller, J., Leichenauer, S., & Wall, A. C. (2016). Proof of the quantum null energy condition. *Physical Review D*, 93, 024017 (arXiv:1509.02542).

- Bousso, R., Fisher, Z., Leichenauer, S., & Wall, A. C. (2016). Quantum focusing conjecture. *Physical Review D*, 93, 064044 (arXiv:1506.02669).
- Bousso, R. (2017). Universal limit on communication. *Physical Review Letters*, 119, 140501 (arXiv:1611.05821).
- Cai, R.-G., & Kim, S. P. (2005). First law of thermodynamics and Friedmann equations of Friedmann-Robertson-Walker universe. *Journal of High Energy Physics*, 0502, 050 (arXiv:hep-th/0501055).
- Casini, H. (2008). Relative entropy and the Bousso bound. *Classical and Quantum Gravity*, 25, 205021 (arXiv:0804.2182).
- Chandrasekaran, V., Longo, R., Penington, G., & Witten, E. (2023). An algebra of observables for de Sitter space. *Journal of High Energy Physics*, 2023(2), 082 (arXiv:2206.10780).
- Chandrasekaran, V., Penington, G., & Witten, E. (2023). Large N algebras and generalized entropy. *Journal of High Energy Physics*, 2023(4), 009 (arXiv:2209.10454).
- Chronis, N., & Sifakis, N. (2025). Conditional cosmological recurrence in finite Hilbert spaces and holographic bounds within causal patches. *Universe*, 12(1), 10.
- Cover, T. M., & Thomas, J. A. (2006). *Elements of Information Theory* (2nd ed.). Wiley.
- Flanagan, É. É., Marolf, D., & Wald, R. M. (2000). Proof of classical versions of the Bousso entropy bound and of the generalized second law. *Physical Review D*, 62, 084035 (arXiv:hep-th/9908070).
- Faulkner, T., & Speranza, A. J. (2024). Gravitational algebras and the generalized second law. *Journal of High Energy Physics*, 2024(11), 099 (arXiv:2405.00847).
- Fischler, W. (2000a, 2000b). Taking de Sitter seriously. Talk given at Role of Scaling Laws in Physics and Biology (Santa Fe); characterized in Bousso (2002).
- Gibbons, G. W., & Hawking, S. W. (1977). Cosmological event horizons, thermodynamics, and particle creation. *Physical Review D*, 15, 2738–2751.
- Hawking, S. W. (1975). Particle creation by black holes. *Communications in Mathematical Physics*, 43, 199–220.
- Hollands, S. (2021). Trace- and improved data processing inequalities for von Neumann algebras. arXiv:2102.07479.
- Jensen, K., Sorce, J., & Speranza, A. J. (2023). Generalized entropy for general subregions in quantum gravity. *Journal of High Energy Physics*, 2023(12), 020 (arXiv:2306.01837).
- Lloyd, S. (2002). Computational capacity of the universe. *Physical Review Letters*, 88, 237901.
- Materov, S. (2026a). The Embedded-Observer No-Go Theorem: Non-Reconstructibility, Non-Autonomy, and Their Quantitative Form for Tensor-Factor Observers in Finite-Dimensional Closed Quantum Systems. DOI 10.5281/zenodo.23066381.
- Materov, S. (2026). GitHub Numerical Companion. https://github.com/SergejMaterov/Principle_of_Significance
- Nielsen, M. A., & Chuang, I. L. (2000). *Quantum Computation and Quantum Information*. Cambridge University Press.
- Podsiadkowska, H., & Wieczorek, R. (2017). Holevo-type bounds for general quantum system. *Reports on Mathematical Physics*, 80(3), 349–360.
- Satishchandran, G., & Sorce, J. (2026). Uniqueness of null-local modular flow. arXiv:2607.20678.
- Segal, I. E. (1960). A note on the concept of entropy. *Journal of Mathematics and Mechanics*, 9(4), 623–629.
- '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, 6377–6396.
- Uhlmann, A. (1977). Relative entropy and the Wigner-Yanase-Dyson-Lieb concavity in an interpolation theory. *Communications in Mathematical Physics*, 54(1), 21–32.
- Van Acoleyen, K. (2010). An operational view on the holographic information bound. *Physical Review D*, 82, 124028 (arXiv:1009.5385).
- Wang, L., & Renner, R. (2012). One-shot classical-quantum capacity and hypothesis testing. *Physical Review Letters*, 108, 200501 (arXiv:1007.5456).

Witten, E. (2022). Gravity and the crossed product. *Journal of High Energy Physics*, 2022(10), 008 (arXiv:2112.12828).

Witten, E. (2024). A background-independent algebra in quantum gravity. *Journal of High Energy Physics*, 2024(3), 077 (arXiv:2308.03663).