

Ordered Patch Theory

Appendix P-3: Fano-Bounded Asymmetric Holography

Anders Jarevåg

April 3, 2026 | DOI: 10.5281/zenodo.19300777

Appendix P-3: Fano-Bounded Asymmetric Holography

Original Task P-3: Fano-Bounded Asymmetric Holography

Problem: Establishing the directionality arrow of holographic equivalence utilizing Fano's Inequality under rate-distortion. **Deliverable:** Formal derivation of the asymmetry.

1. Introduction: The Tension with Exact Duality

Standard formulations of the holographic principle (e.g., AdS/CFT duality) posit an *exact isomorphism* between a higher-dimensional bulk and its lower-dimensional boundary. In pure quantum gravity formulations, these descriptions are mathematically perfectly symmetric; the state in the bulk uniquely specifies the state on the boundary, and crucially, the state on the boundary uniquely specifies the state in the bulk. Neither representation is ontologically prior.

The Ordered Patch Theory (OPT) structurally disrupts this symmetry. OPT asserts that the algorithmic generative rules (\mathcal{I} , "Substrate") hold ontological priority, while the phenomenological world (R , "Render") is a derived predictive shadow. This asymmetry introduces a formal theoretical tension: if holographic dualities arise organically from informational coding limits, why does the symmetry strictly break in our local causal patch?

This appendix resolves the tension by deploying **Fano's Inequality** under the Solomonoff algorithmic measure. We formally derive (conditional on Assumption P-3.1, established in §2) that the structural requirement of a phenomenological observer inherently transforms holography from a symmetric duality into an **Asymmetric One-Way Holographic** projection.

2. The Stability Filter as a Lossy Compression Map

In OPT, the phenomenological world exists solely within the narrow bandwidth geometry of the conscious integration channel. The foundational algorithm \mathcal{I}

operates over a Solomonoff universal semimeasure environment.

We must formally establish that the transformation to the observer target render R via the Stability Filter (Φ) is a fundamentally lossy map. To do so without circular reasoning, we deploy the defining Markov Chain sequence separating the substrate from the render via the codec’s Markov boundary $X_{\partial A}$ (by the Markov blanket separability condition, Preprint §3.4 / Eq. 8):

$$\mathcal{I} \rightarrow X_{\partial A} \rightarrow R$$

By the Data Processing Inequality, information cannot increase through successive transformations. Therefore, the mutual information strictly scales as:

$$I_m(\mathcal{I}; R) \leq I_m(X_{\partial A}; R)$$

Note: The mutual information I_m is formally defined here under a normalised version of the Solomonoff semimeasure ($p(\nu) = m(\nu) / \sum_{\nu \leq K_{\max}} m(\nu)$) over the finitely many algorithms bounded below the complexity limit K_{\max} .

Because the Stability Filter imposes a channel capacity C_{\max} on the mapping to the boundary ($X_{\partial A} \rightarrow R$), Shannon’s fundamental channel capacity theorem dictates that $I_m(X_{\partial A}; R) \leq T \cdot C_{\max}$ holds for any input distribution, including the normalized Solomonoff prior. Chaining this inequality with the DPI establishes that the phenomenological render is strictly bounded by a finite bottleneck integrated over duration T . Therefore:

$$I_m(\mathcal{I}; R) \leq T \cdot C_{\max}$$

For exact symmetric duality to hold, the mapping between the bulk (Substrate) and the boundary (Render) must be perfectly invertible ($\Phi^{-1} : R \rightarrow \mathcal{I}$). Let ν_{true} represent the specific, unknown realization of the generating algorithm responsible for our observed universe (drawn from the underlying algorithmic ensemble). Let N effectively represent the vast combinatorial space of lower-semicomputable algorithms operating beneath a finite complexity threshold constraint K_{\max} .

Assumption P-3.1 (Substrate Complexity Scaling): The true generating complexity satisfies $K(\nu_{\text{true}}) \gg T \cdot C_{\max}$. (This is explicitly motivated by the Minimum Description Length parsimony arguments in Appendix T-4; any algorithm encoding equivalent Standard Model physics requires immense structural data).

Because the mutual information bound falls severely short of the information required to specify the true bulk state under Assumption P-3.1, Φ is strongly established as a **lossy compression map**.

3. Conditional Entropy and the Conditioned Solomonoff Prior

To quantify the cost of this lossy compression, we evaluate the error probability of an observer O situated strictly inside the render R attempting to uniquely infer the true underlying generative substrate algorithm (ν_{true}).

Crucially, evaluating expected complexity over the raw Solomonoff prior creates a paradox: the raw prior is heavily dominated by the low- K tail (trivial programs and constant sequences). Over the unconditioned universal distribution, the expected complexity $\langle K \rangle_M$ evaluates to a tiny negligible boundary ($O(100)$ bits). If this held true for our universe, the condition $K(\nu_{\text{true}}) \gg T \cdot C_{\text{max}}$ would instantly fail, collapsing the entropic boundary entirely.

However, the raw unconditioned prior is structurally meaningless for generating internal phenomenological observers. To possess sufficient physical mechanics, “requisite variety,” and temporal duration to host an active-inference self-model, the generating algorithm must possess an immense minimal structural baseline. As quantified in Appendix T-4 §2.1, the full generating algorithm ν_{true} must encode not only the Standard Model law structure ($K(\text{laws}) \approx 1750$ bits) but also the specific microstate initial conditions, which by Penrose’s estimate requires $K(\text{IC}|\mathcal{M}_1) \sim 10^{123}$ bits. The combined complexity $K(\nu_{\text{true}}) \sim 10^{123}$ bits therefore sets $K_{\text{threshold}}$. We must therefore evaluate entropy exclusively over the **Stability-Filter-Conditioned Prior** ($M|SF$)—the subset of generating algorithms possessing sufficient complexity ($K(\nu) \geq K_{\text{threshold}} \sim 10^{123}$) to generate a universe consistent with the specific observed physics and initial conditions of this causal patch. Confirming $\langle K \rangle_{M|SF} \geq K_{\text{threshold}} \gg T \cdot C_{\text{max}}$, we independently anchor Assumption P-3.1 and confirm over this conditioned distribution that the expected generative complexity correctly bounds to $\langle K \rangle_{M|SF} \gg T \cdot C_{\text{max}}$.

The true structural conclusion of informational starvation operates decisively through Shannon Conditional Entropy within this restricted parameter space:

$$H_{m|SF}(\mathcal{I}|R) = H_{m|SF}(\mathcal{I}) - I_{m|SF}(\mathcal{I}; R) \approx \langle K \rangle_{M|SF} - T \cdot C_{\text{max}} \approx \langle K \rangle_{M|SF}$$

(Note: The channel capacity bound $T \cdot C_{\text{max}}$ applies universally to $I_{m|SF}$ by the exact same Shannon supremum theorem established in Section 2 for any input distribution.)

(Note: The identity $H_m(\mathcal{I}) \approx \langle K \rangle_M$ is precise save for a negative normalization constant $\log_2 Z$ natively inherited from the truncated Solomonoff semimeasure limit. Because the full unrestricted semimeasure approaches 1, $|\log_2 Z|$ is safely bounded directly by the description-length overhead of the Universal Turing Machine, a fixed constant $c_U \approx O(100)$ bits. This establishes the structural normalization as a trivial rounding error against the functionally macroscopic scale of $\langle K \rangle_{M|SF}$.)

Corollary: Kolmogorov-Weighted Fano Inequality

While the conditional entropy bound overwhelmingly serves as the physical proof of macroscopic informational starvation, we observe that adapting Fano's Inequality under the same normalized, SF-conditioned Solomonoff-weighted measure replaces the standard uniform entropy term in the numerator with the expected Kolmogorov Complexity, generating the secondary statistical lower bound:

$$P(\hat{\mathcal{I}} \neq \mathcal{I}) \geq \frac{\langle K \rangle_{M|SF} - T \cdot C_{\max} - 1}{K_{\max}}$$

(Note: Under Assumption P-3.1, this error probability floor is strictly positive but mathematically weak relative to macroscopic scales; it should be understood solely as a secondary statistical corollary to the entropic limit.)

4. Connection to QECC Constraints and Informational Irreversibility

Because the conscious integration bound $T \cdot C_{\max}$ evaluates to a negligible fraction in comparison to the algorithmic source, the conditional entropy $H_{m|SF}(\mathcal{I}|R)$ remains approximately identical to $\langle K \rangle_{M|SF}$. The vast majority of the generative substrate information is irreducibly inaccessible from within R .

CONJECTURE (Open Edge): The conditional entropy limits defined here map cleanly onto the breakdown defined in Appendix P-2. Under the Quantum Error Correction Code (QECC) bulk-boundary map, the Stability Filter Φ acts as the partial isometry protecting the low-energy boundary states. We conjecture that the specific MERA cut depth τ^* relates mathematically to the precise spatial capacity horizon defining the conditional entropy starvation threshold, bounded by the algorithmic QECC ADH reconstruction condition. A formal derivation bridging the statistical limits and the geometric MERA cut depth is deferred to future theoretical work. Nonetheless, functionally, information located deeper in the bulk substrate is permanently encrypted by the one-way compression.

The observer's internal reconstruction effort is mathematically saturated. The inverse map Φ^{-1} is **statistically non-invertible** on the minimum-error scale from within R .

5. Conclusion: Phenomenological Priority

Therefore, the informational arrow operates dominantly in one direction: information is systematically destroyed during the projection from Substrate to Render, and it cannot be causally or statistically recovered from within the phenomenological frame.

Through this Fano boundary formulation, under Assumption P-3.1, we formally establish that **Asymmetric Holography** is a strict mathematical consequence of placing a rate-limited observer inside the causal framework.

- The Substrate \mathcal{I} is the fundamental engine because its total state fully determines the conditional probability distribution mapped by Φ .
- The Render \mathcal{R} is strictly secondary because its state is a reduced summary incapable of retroactively predicting the substrate causing it.

Phenomenological consciousness is thus the first-person internal experiential state of being structurally trapped on the output side of a non-invertible compression algorithm. It establishes that while our local physics obeys holographic constraints (optimizing area vs. volume limits), the boundary representation acts as an irreversible epistemic bottleneck, formally breaking the symmetry required for standard exact string-theoretic duality.