

Objectivity from Records: An Exponential Consensus Bound

A Technical Note on the S_I Signature

Jeremy C. Jones

HoldingLight LLC — ORCID 0009-0007-2515-3774 — universalcollapse.com

Series: Universal Collapse Theory — Technical Note • Version v1.0 • 2026-05 • CC BY 4.0

Abstract

Principle. If macroscopic regularities arise from constraint-guided collapse and become objective through redundant records, then independent informative records should force inter-observer agreement at a quantifiable rate. **Result.** We formalize the simplest case as a binary hypothesis test: a realized outcome X is redundantly recorded in conditionally independent fragments $Y_1, \dots, Y_k$. Under positive Chernoff information per fragment, the optimal estimator’s error — and hence pairwise observer disagreement — decays exponentially in k . This yields a cross-domain agreement curve that is testable in any setting with redundant readouts. Heterogeneous, correlated, and non-binary cases are noted as remarks rather than formalized. The result is conditional on the stated independence and discriminability assumptions; it is operational, not philosophical.

Review target

This technical note does not ask the reader to accept Universal Collapse Theory as a metaphysical system or to accept a complete theory of objectivity. It asks whether one core S_I claim can be formalized: if a realized outcome is redundantly encoded in conditionally independent, discriminable record fragments, then optimal observers using those fragments converge exponentially fast. The note should be accepted only as a conditional bound. It should be revised or rejected if the independence or discriminability assumptions are misstated, if the exponential bound is mathematically incorrect, or if the empirical agreement-curve protocol fails to distinguish genuine independent redundancy from correlated pseudo-redundancy.

Stack placement

This note is a formal support document for the Records–SoE–UIS layer of the Universal Collapse Theory library. *Records Across Nature, Life, and Mind* (Jones, 2026a) defines records as the persistence layer and states S_I (“redundancy $\rightarrow$ consensus”) as a portable empirical signature. *The Structuralization of Empiricism* (Jones, 2026b) uses S_I to describe how redundant records stabilize objectivity and empirical convergence. The *Update Integrity Standard* (UIS v1.0; Jones, 2026c) supplies the independence-audit and k_{eff} reporting discipline. The present note proves the simplest binary case: independent informative records imply an exponential consensus curve. Operational pair: Methods- S_I (*Auditing Independence in Multi-Channel Measurement*). Together they constitute the S_I formal–protocol pair — this Technical Note proves the formal bound; the paired Methods Paper translates it into a deployable audit protocol. It is a citable lemma for downstream S_I empirical demonstrations, not a standard.

1. Setup and definitions

Outcome. A realized binary variable $X \in \{0, 1\}$ with equal priors. Generalizations to finite $|X| > 2$ are noted in §3.

Fragments. Each observer reads k records $Y_1, \dots, Y_k$ taking values in a measurable space. Conditional on $X = x$, each Y_i has density p_x with respect to a common dominating measure μ .

Assumption A (conditional independence). $Y_1, \dots, Y_k$ are conditionally independent given X .

Assumption B (discriminability). The pair $\{p_0, p_1\}$ has Chernoff information at least $c_0 > 0$:

$$C(p_0, p_1) := \sup_{0 \leq s \leq 1} [-\log \int p_0(y)^s p_1(y)^{1-s} d\mu(y)] \geq c_0 > 0.$$

Equivalently, $C(p_0, p_1) = -\log \inf$ of the integral over $s \in [0, 1]$. At the endpoints $s = 0$ and $s = 1$ the integral equals 1, contributing zero to the negated log; the supremum is therefore achieved on the interior, where the integrand can be strictly less than 1.

Estimator. Each observer applies the Bayes-optimal (MAP) rule on its k -fragment sample.

Disagreement. Two observers read disjoint conditionally independent k -fragment samples and report $\hat{X}_k^{(1)}, \hat{X}_k^{(2)}$.

2. Theorem 1 (iid binary case)

Theorem 1 (Redundancy $\Rightarrow$ Exponential Consensus). Let $X \in \{0, 1\}$ with equal priors. Conditional on $X = x$, let $Y_1, \dots, Y_k$ be iid with density p_x . Suppose $C(p_0, p_1) \geq c_0 > 0$ (Assumption B). Then:

(a) The MAP error satisfies $P_e(k) = \Pr[\hat{X}_k \neq X] \leq \frac{1}{2} e^{-kc_0}$.

(b) For two observers reading disjoint conditionally independent k -fragment samples, pairwise disagreement satisfies

$$\Pr[\hat{X}_k^{(1)} \neq \hat{X}_k^{(2)}] \leq e^{-kc_0}.$$

(c) **Sample complexity.** To drive pairwise disagreement below $\delta \in (0, 1)$, it suffices that $k \geq (1/c_0) \log(1/\delta)$.

Proof.

The log-likelihood ratio $A_k = \sum_i \log[p_1(Y_i) / p_0(Y_i)]$ is a sum of k iid random variables under either hypothesis. Chernoff's bound for binary hypothesis testing with equal priors gives $P_e(k) \leq \frac{1}{2} \exp(-kC)$, where $C = C(p_0, p_1) \geq c_0$ by Assumption B (Chernoff 1952; Cover & Thomas 2006, Ch. 11). The pairwise-disagreement bound follows from a union bound over the two independent MAP decisions:

$$\Pr[\hat{X}_k^{(1)} \neq \hat{X}_k^{(2)}] \leq \Pr[\hat{X}_k^{(1)} \neq X] + \Pr[\hat{X}_k^{(2)} \neq X] \leq 2 P_e(k) \leq e^{-kc_0}. \blacksquare$$

3. Remarks and extensions

Heterogeneous fragments (Bhattacharyya form). If fragment i has its own density $p_{i,x}$, choosing $s = \frac{1}{2}$ gives the looser but additive Bhattacharyya bound

$$P_e(k) \leq \frac{1}{2} \exp(-\sum_i B_i), \quad \text{where} \quad B_i = -\log \int \sqrt{p_{i,0}(y) p_{i,1}(y)} d\mu_i(y).$$

If each $B_i \geq b_0 > 0$, exponential decay in k is recovered with rate b_0 . A tighter heterogeneous bound uses a common s optimized across fragments; we omit the formal statement.

Beyond binary outcomes. For finite $|X| > 2$, applying the bound to each pair of classes gives exponential decay at the worst pairwise rate. The empirical signature is unchanged in form.

Correlated fragments. Under some weak-dependence regimes (e.g., summable α -mixing), an effective sample size $k_{\text{eff}} \leq k$ can often model the degradation from correlation, and exponential-like decay may persist over the audited range. A rigorous bound requires domain-specific dependence assumptions; in practice, k_{eff} should be audited per UIS protocols, not assumed. This is the boundary at which “independent redundancy” must be distinguished from “correlated pseudo-redundancy.”

Quantum readouts. When records are obtained as classical readouts from environmental fragments of a quantum system (the Quantum Darwinism setting; Zurek 2009), the present theorem applies to the readouts. A fully quantum version would replace the classical Chernoff information with the quantum Chernoff bound; we do not pursue that here.

4. Empirical signature (agreement curve)

Prediction. In any system where a single latent outcome X is redundantly recorded in conditionally independent fragments, pairwise observer disagreement decays approximately as ae^{-bk} .

Protocol. Vary k . Report the independence audit and k_{eff} per UIS conventions. Fit $\Pr[\text{disagree}]$ vs k (or k_{eff}) to ae^{-bk} ; report the exponent b with confidence intervals.

Failure condition. S_I fails locally if disagreement does not decay with independently audited k_{eff} , or if the curve saturates earlier than the audited k_{eff} predicts. The latter is the operational diagnostic for correlated fragments masquerading as independent.

Domains. Multi-detector physics readouts; quantum-Darwinism-style experiments; ensemble sensors with disjoint sample shards; human annotation tasks with non-overlapping evidence shards.




Figure 1. Agreement curve from a binary symmetric channel simulation. A latent outcome $X \in \{0, 1\}$ has equal priors. Each fragment Y_i is a noisy bit with crossover probability p , conditionally independent given X . Two observers each read a disjoint set of k fragments and apply MAP (= majority vote for this channel). Points are empirical pairwise disagreement rates over 20,000 trials per k ; dashed lines are the theoretical Chernoff bound e^{-kC} with $C = -\log[2\sqrt{p(1-p)}]$. The finite-sample bound is loose for small k , while the Chernoff exponent captures the asymptotic decay rate as k grows. Reproduction code in Appendix A.

5. Minimal framework hook

We model realization as collapse under constraints: a constraint set K selects admissible outcomes from Ω via a map C^K . Records R are durable, informative, non-destructively readable traces that can constrain later updates. Redundancy means many conditionally independent records of the same X . Theorem 1 shows that redundancy forces consensus at a quantifiable exponential rate — an operational bridge from records to objectivity.

This is *operational* objectivity: agreement between optimal observers on a shared latent outcome, given independent redundant records of it. It is not a claim that observers cannot disagree, that records cannot be corrupted, or that all redundancy is genuine. The independence audit — not the bound itself — is what does the work in any real domain.

6. Limitations

The result is conditional on Assumptions A and B. Independence is the demanding assumption; discriminability is usually the easier one. The bound is operational, not philosophical: it characterizes inter-observer agreement under MAP decisions, not “objectivity” in any broader sense. For quantum systems, the theorem applies to classical readouts; the fully quantum case requires the quantum Chernoff bound. Heterogeneous and correlated extensions are noted but not formalized; rigorous treatment is left to domain-specific work and to the future T16 empirical demonstration that this note supports.

Appendix A. Reproduction code

The following Python script reproduces Figure 1 (binary symmetric channel; equal priors; for the plotted values $p < 1/2$, MAP reduces to majority vote with random tie-break for even k ; 20,000 trials per k).

Dependencies: NumPy and Matplotlib only.

```
import numpy as np
import matplotlib.pyplot as plt

def disagreement_rate(k, p, n_trials=20000, seed=0):
    rng = np.random.default_rng(seed + k * 7919)
    X = rng.integers(0, 2, size=n_trials)
    flips1 = (rng.random((n_trials, k)) < p).astype(int)
    flips2 = (rng.random((n_trials, k)) < p).astype(int)
    Y1 = X[:, None] ^ flips1
    Y2 = X[:, None] ^ flips2
    half = k / 2.0
    Xhat1 = (Y1.sum(axis=1) > half).astype(int)
    Xhat2 = (Y2.sum(axis=1) > half).astype(int)
    # random tie-break for even k
    tie1 = (Y1.sum(axis=1) == half)
    tie2 = (Y2.sum(axis=1) == half)
    if tie1.any(): Xhat1[tie1] = rng.integers(0, 2, size=tie1.sum())
    if tie2.any(): Xhat2[tie2] = rng.integers(0, 2, size=tie2.sum())
    return (Xhat1 != Xhat2).mean()

def C_bsc(p):
    # Chernoff information for the binary symmetric channel:
    # C = -log[ 2 * sqrt(p*(1-p)) ]
    return -np.log(2 * np.sqrt(p * (1 - p)))

ks = list(range(1, 81, 2))
for p in [0.45, 0.40, 0.35]:
    rates = [disagreement_rate(k, p) for k in ks]
    plt.semilogy(ks, rates, "o", label=f"p={p}")
```

```
C = C_bsc(p)
plt.semilogy(ks, np.exp(-np.array(ks) * C), "--", alpha=0.5)
plt.xlabel("k"); plt.ylabel("Pr[disagree]"); plt.legend(); plt.show()
```

References

- Chernoff, H. (1952). A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *The Annals of Mathematical Statistics*, 23(4), 493–507.
<https://doi.org/10.1214/aoms/1177729330>
- Cover, T. M., & Thomas, J. A. (2006). *Elements of Information Theory* (2nd ed.). Wiley.
- Jones, J. C. (2026a). *Records Across Nature, Life, and Mind*. HoldingLight LLC.
<https://doi.org/10.17605/OSF.IO/7H6DY>
- Jones, J. C. (2026b). *The Structuralization of Empiricism*. HoldingLight LLC.
<https://doi.org/10.17605/OSF.IO/J4GZ9>
- Jones, J. C. (2026c). *Update Integrity Standard (UIS v1.0)*. HoldingLight LLC.
<https://doi.org/10.17605/OSF.IO/DWM29>
- Zurek, W. H. (2009). Quantum Darwinism. *Nature Physics*, 5, 181–188.
<https://doi.org/10.1038/nphys1202>
- This paper is part of the Universal Collapse Theory library. For a reading guide and full architecture, visit universalcollapse.com/roadmap.

AI Disclosure

AI tools were used to assist with manuscript preparation. The underlying theory, arguments, and interpretive claims are the author's own, and the author takes full responsibility for the content.

Citation

Jones, J. C. (2026). *Objectivity from Records: An Exponential Consensus Bound* (UCT Technical Note v1.0). HoldingLight LLC.