

THE ZERO PARADOX

ZP-C: Information Theory

Version 1.14 | May 2026

This document is self-contained within information theory and discrete analysis on Q_2 . The topological structure of Q_2 — specifically total disconnectedness (ZP-B T5), the clopen ball hierarchy, and the binary existence axiom (AX-B1) — is imported from ZP-B as a dependency. Every claim is marked as Derived, Axiomatic, Defined, or Candidate.

Encoding commitment. The two ontological states of AX-B1 are represented here as point-mass (Dirac) distributions over $\{0, 1\}$ — the minimum-sufficient probabilistic encoding (RP-1, Section II). Under this encoding, the information-theoretic separation between the null state and the first atomic state is exactly 1 bit (T1b). The 1-bit result is conditional on this encoding; RP-1 declares and justifies the commitment. Separately, at \perp , the primary descriptive tool of information theory — surprisal — is unbounded above (L-INF, Section III): no finite external description can contain the null state. These two results — the 1-bit cost of the transition and the infinite descriptive cost of the origin — are the information-theoretic expression of the same structural constraint on the null state.

Illustrated Companion: A paired ZP-C Illustrated Companion provides concrete examples and visual intuitions for the results here. Examples are kept separate from the formal layers to distinguish illustrative material from proofs.

I. Kolmogorov Complexity and the Incompressibility Threshold

Let x be a binary string of length n . The conditional Kolmogorov complexity is:

$$K(x|n) = \min \{ |p| : U(p, n) = x \}$$

where U is a fixed universal Turing machine and $|p|$ is the length of program p .

Definition D1 — Incompressibility Threshold

$$P_0 = \lim_{n \rightarrow \infty} K(x|n) / n$$

P_0 is the algorithmic entropy rate of x . When P_0 reaches its maximum, x has exhausted its capacity for self-description.

Remark R1 — Scope of P_0

What P_0 establishes (Derived): the point at which x becomes incompressible.

What P_0 does not establish (prior versions: Axiomatic): that incompressibility causes the Binary Snap. Version 1.4 updates this: the generative claim is now a Candidate Theorem (T-BUF), not a bare axiom. See Section V.

II. Informational Work: State Representations and JSD

2.1 The Representation Principle

Principle RP-1 — Minimum Sufficient Probabilistic Representation

The minimum sufficient probabilistic representation of a binary ontological state is a point-mass (Dirac) distribution over $\{0,1\}$: all probability mass is assigned to the value the state occupies, and none to the other value.

Justification: A point-mass distribution is the unique distribution that is (a) faithful — assigns non-zero probability only to the ontologically actual value, and (b) non-redundant — carries no uncertainty about which state is occupied. Any distribution with mass at both values represents partial existence and partial non-existence simultaneously. AX-B1 excludes this.

Status: PRINCIPLE — representational commitment parallel to MP-1 in ZP-B.

Theorem T1 — State Representations are Uniquely Derived

Given AX-B1 and RP-1, the distributions are unique:

$P = (1, 0)$ (Null State: all mass at 0 — non-existence)

$Q = (0, 1)$ (First Atomic State: all mass at 1 — existence)

Proof: AX-B1 establishes two ontological states. RP-1 requires point-mass representation of each. The Null State occupies value 0: $P = (1,0)$. The First Atomic State occupies value 1: $Q = (0,1)$. No distribution $(p, 1-p)$ with $0 < p < 1$ is consistent with AX-B1. P and Q are unique. Status: DERIVED from AX-B1 and RP-1. ✓

Corollary T1b — JSD = 1 bit

For $P=(1,0)$ and $Q=(0,1)$ with $M=(1/2, 1/2)$:

$D_{KL}(P||M) = 1$ bit, $D_{KL}(Q||M) = 1$ bit, $JSD(P||Q) = 1$ bit

$E = JSD(P||Q) = 1$ bit [information-theoretic, dimensionless]

Status: Derived from AX-B1 and RP-1 via T1. E is not physical energy in joules. Dimensional bridge belongs to ZP-E as BA-1.

Definition D3 — Dirac Measure δ_0

$$\int_{\Omega} f d\delta_0 = f(0)$$

δ_0 places unit mass at 0. Compatible with the discrete topology of $\{0,1\}$ and with the totally disconnected topology of Q_2 (ZP-B T5). δ_0 governs behaviour exactly at $x = 0$.

Remark R2 — Hamming Cross-Validation

$$d_H(P, Q) = 1 \text{ [Hamming, on } \{0,1\}]$$

The agreement between $d_H = 1$ and $E = 1$ bit is a consistency check, not a proof that Hamming distance and JSD are equivalent in general. Both are computed by independent methods on the same AX-B1-derived distributions.

III. The Discrete Surprisal Field on Q_2

Remark R3 — Why the Smooth Embedding Remains Retired

ZP-B T5 establishes Q_2 is totally disconnected. ZP-B T2 establishes every ball is clopen. Importing smooth calculus operators onto a totally disconnected space imports a smoothness assumption that contradicts ZP-B. The smooth embedding, MO-1, and P1 from v1.1 remain retired.

Definition D4 — Discrete Surprisal Function I

For $x \in Q_2$ with $x \neq 0$ and $P(x) > 0$:

$$I(x) = -\log_2 P(x)$$

As $v_2(x) \rightarrow +\infty$ (x approaches 0 in the 2-adic metric), $I(x) \rightarrow +\infty$: states 2-adically close to 0 are informationally extreme.

The probability measure P on $Q_2 \setminus \{0\}$ is the branching measure induced by the binary ball hierarchy of Q_2 (RP-2 below): at each level k , the two sub-balls of $B(0, 2^{-k})$ each receive half the probability mass of their parent ball.

Principle RP-2 — Branching Measure on $Q_2 \setminus \{0\}$

The probability measure P on $Q_2 \setminus \{0\}$ is the canonical branching measure induced by the binary ball hierarchy: at each level k , each sub-ball receives half the probability mass of its parent. This is a representational commitment — a specific measure choice not uniquely forced by AX-B1 alone.

Justification: The branching measure assigns equal weight to each binary branch, consistent with AX-B1's symmetric treatment of existence and non-existence. Alternative measures (e.g. non-uniform branching weights) would depart from this symmetry without additional motivation.

Parallel: RP-1 fixes the point-mass representation of ontological states. RP-2 fixes the probability measure on the 2-adic domain. Both are representational commitments sitting alongside, not derived from, AX-B1.

Status: PRINCIPLE — representational commitment. Required by T2 and L-INF.

Definition D5 — Discrete Surprisal Difference Operator DF

For $x, y \in Q_2 \setminus \{0\}$:

$$DF(x, y) = I(y) - I(x) = \log_2 [P(x) / P(y)]$$

DF is antisymmetric: $DF(x, y) = -DF(y, x)$. No smoothness is assumed; DF is defined entirely pointwise.

Definition D6 — Discrete Circulation (Extended)

FINITE CASE: A discrete loop $\gamma_n = (x_0, x_1, \dots, x_n, x_0)$ in $Q_2 \setminus \{0\}$ returning to its start. Circulation: $C(DF, \gamma_n) = \sum_{i=0}^n DF(x_i, x_{i+1})$ where $x_{n+1} = x_0$.

Note: By telescoping, $C(DF, \gamma_n) = I(x_0) - I(x_0) = 0$ for all finite loops. DF is conservative on all finite loops.

Definition D6 — Discrete Circulation (Extended)

INFINITE CASE: An infinite sequence $\sigma = (x_1, x_2, \dots)$ in $\mathbb{Q}_2 \setminus \{0\}$ where $v_2(x_i) = i$ for all $i \geq 1$. Partial circulation: $C_n(\text{DF}, \sigma) = I(x_{n+1}) - I(x_1)$.

Circulation of σ : $C(\text{DF}, \sigma) = \lim_{n \rightarrow \infty} C_n(\text{DF}, \sigma)$ if the limit exists or diverges.

An infinite sequence σ exhibits non-conservative behaviour if $C(\text{DF}, \sigma)$ diverges.

Theorem T2 — DF Exhibits Non-Conservative Behaviour on Infinite Sequences Approaching 0

DF is conservative on all finite loops ($C = 0$ by telescoping). On infinite sequences through the ball hierarchy approaching 0, the circulation diverges: $C(\text{DF}, \sigma) \rightarrow +\infty$.

Proof: Let $\sigma = (x_1, x_2, \dots)$ with $x_i = 2^{-i}$. Then $P(x_i) = 2^{-i}$, so $I(x_i) = i$. Partial circulation $C_n = I(x_{n+1}) - I(x_1) = n$. As $n \rightarrow \infty$: $C(\text{DF}, \sigma) = +\infty$.

Divergence arises from the unbounded depth of the 2-adic ball hierarchy (ZP-B). Finite conservation and infinite divergence are distinct regimes; no contradiction arises.

Status: DERIVED from ZP-B ball hierarchy structure and branching measure on \mathbb{Q}_2 . OQ-C1 closed. ✓

III-B. Informational Extremity of the Null State

Lemma L-INF — Informational Extremity of \perp

The surprisal $I(n) = n$ at ball-hierarchy depth n is unbounded above: for any finite bound M , there exist depths n with $I(n) > M$.

The null state $\perp = c_0$ corresponds to the limit point $0 \in \mathbb{Q}_2$ — the limit of the binary ball hierarchy at infinite depth. The binary branching measure assigns equal probability mass at each branch level (D4), and the surprisal diverges without bound as depth increases (T2). At this limit, no finite bound M contains the informational content of \perp .

Proof: Let $M \in \mathbb{R}$. By the Archimedean property, $\exists n \in \mathbb{N}$ with $n > M$. Then $I(n) = n > M$. Since M was arbitrary, surprisal is unbounded above. ✓

Formal content: surprisal is not bounded above by any real M .

Semantic content: \perp is informationally extreme — it is the limit point of the binary ball hierarchy under the 2-adic metric, the accumulation point approached by sequences of increasing depth. No finite bound M contains the surprisal at that limit; therefore no finite external interpreter can hold \perp as a static description. This is the mathematical premise for DA-1 (ZP-E § I-DA1). ZP-A CC-2 ($\perp = \{\perp\}$) provides a structural second grounding for the same conclusion: a self-containing object has no external interpreter by structure (ZP-A R3). The informational argument from the ball hierarchy and the structural argument from self-containment are independent derivations converging on the same fact.

Note: the connection from informational extremity to forced execution is a named design principle (DA-1 in ZP-E), not a mathematical consequence of L-INF alone. L-INF supplies the formal premise; DA-1 supplies the ontological bridge.

Lemma L-INF — Informational Extremity of \perp

Status: DERIVED from D4 and T2. Structural corroboration: ZP-A CC-2 ($\perp = \{\perp\}$) and R3. Lean: ZPC.l_inf (purity check: no non-Mathlib axioms).

Remark R-BRIDGE — On the Relationship Between K and 2-Adic Surprisal

ZP-C uses two complexity measures: Kolmogorov complexity $K(x|n)/|x|$ (Section I) and 2-adic surprisal $I(n) = n$ (Section III). These measure related but distinct structures. Their relationship warrants explicit statement.

What each measures. $K(x|n)/|x| = 1$ is a property of a specific binary string x : no program shorter than x generates it. It is a pointwise statement about a particular configuration at the incompressibility threshold. $I(n) = n$ is a property of position in the ball hierarchy: the surprisal of any string at 2-adic depth n is n , by the branching measure RP-2. It is a structural property of depth, not of any particular string.

Where they coincide. At P_0 , both conditions hold simultaneously for the same configuration: $K(c_1|n)/|c_1| = 1$ (no shorter external description exists) and $I(n) = n \rightarrow \infty$ (surprisal grows without bound as depth increases). This is not a coincidence of definition. The 2-adic limit point $0 \in Q_2$ is the unique accumulation point of the binary ball hierarchy; algorithmic incompressibility is the K-complexity characterisation of the same extremality. Both locate the same threshold from independent directions. The structural argument for why this convergence is not accidental — that these two independently defined measures are forced to meet at the same limit by the framework's architecture — is developed in ZP-E Remark R-AFA.

Where they diverge. Away from P_0 , the measures are not interchangeable. $K(2^n) = O(\log n)$: the n -fold power of 2 is compactly described by a short program, yet $v_2(2^n) = n$ grows without bound. A string can be 2-adically deep without being Kolmogorov-incompressible, and vice versa. The two measures agree at the limit but diverge throughout the interior.

How ZP-C uses them. ZP-C uses K and I as independent convergent routes to the same conclusion — that \perp at P_0 admits no finite external interpreter — not as a single unified measure. L-INF is the 2-adic surprisal argument: $I(n)$ unbounded \rightarrow no finite interpreter can hold \perp . DA-1 Path 3 (ZP-E) is the K-incompressibility argument: $K = 1$ eliminates the static-description state via D7's dichotomy. Their independence is preserved deliberately. The convergence of two structurally distinct complexity measures on the same threshold is part of the argument for DA-1, not a circularity within it. Axiomatising an equivalence between K and I would merge these two routes into one and eliminate that independence.

IV. Execution as State: The Hardware Lemma

Definition D7 — Machine Configuration

A machine configuration c is a complete description of a Turing machine at a given moment: the current state of the control unit, the position of the read/write head, and the contents of all tape cells.

c_0 : the initial configuration — the machine exists but has not yet begun execution (no instruction has been fetched).

c_1 : the first running configuration — the machine has fetched and begun executing its first instruction.

Note: D7 does not import physics. It is a structural distinction between two discrete machine states within the standard Turing model.

Lemma L-RUN — Execution is a Non-Null State Change

The transition from c_0 (initial: not yet running) to c_1 (running: first instruction fetched) is a discrete, irreducible state change in the machine configuration.

Proof:

Step 1 — c_0 and c_1 are distinct configurations by D7: the control unit is in a different state (idle vs. executing). $c_0 \neq c_1$.

Step 2 — By AX-B1, a state either exists or it does not. The machine at c_1 occupies a configuration distinct from c_0 . The transition $c_0 \rightarrow c_1$ is a binary state change in the sense of AX-B1.

Step 3 — The transition is irreducible: there is no intermediate configuration between c_0 and c_1 . The first instruction fetch is the minimal unit of execution.

Step 4 — c_0 is modeled as corresponding to \perp in the semilattice (CC-2 — see below). c_1 is strictly above \perp : the machine configuration has gained content (an active execution context) not present in c_0 .

Conclusion: $c_1 \neq \perp$. The act of execution is itself a non-null state, regardless of what the output tape contains. ✓

Status: DERIVED from AX-B1 and D7. No Coding Theorem required. No output-tape contents required.

Conditional Claim CC-2 — $c_0 = \perp$ (Parallel to CC-1 in ZP-A)

We model the initial machine configuration c_0 as corresponding to \perp in the semilattice (L, \vee, \perp) . This is not derived from D7 or from A1–A4 — it is a modeling choice.

Motivation: c_0 is the machine state prior to any execution: no instruction has been fetched, no content has accumulated. \perp is the algebraic element below all others — the state of no accumulated content. The identification is motivated by structural correspondence, not derivation.

Parallel: CC-1 in ZP-A commits to $S_0 = \perp$ for state sequences. CC-2 makes the same commitment for machine configurations. Both are modeling choices, not derivations from the core axioms.

Effect: Step 4 of L-RUN depends on CC-2. The conclusion $c_1 \neq \perp$ follows from L-RUN's proof; the identification of c_0 with \perp is the additional modeling commitment supplied by CC-2.

Status: CONDITIONAL CLAIM — modeling commitment; not derived from D7 or A1–A4.

Remark R4 — Configuration and Output Are Independent

D-TQ-2 (Output Definition): A computation's output is defined by the contents of its tape/register at termination.

L-RUN establishes that configuration state and output state are independent. A program may terminate with null output while still having passed through non-null configuration states during execution.

The question "can a program output \perp without any non-null intermediate state?" is answered by looking at the execution trace, not the output tape. This distinction is load-bearing for T-BUF and TQ-IH.

V. The Test Question and the Buffer Overflow Theorem

Test Question TQ-IH — Can a Program Output \perp Without a Non-Null Intermediate State?

Question: Does there exist a program p such that $U(p) = \perp$ and the execution trace $\tau(p)$ contains no configuration c_i with $c_i \neq \perp$?

Answer: No.

Proof:

Step 1 — Any program that executes must pass through c_1 (by definition of execution, D7).

Step 2 — By L-RUN, $c_1 \neq \perp$.

Step 3 — Therefore the execution trace $\tau(p) = (c_0, c_1, \dots)$ contains c_1 , which is a non-null configuration state, regardless of what appears on the output tape at termination.

Step 4 — The output of p being \perp does not imply that all configurations in $\tau(p)$ are \perp . Output state and configuration state are distinct (R4).

Conclusion: No program can produce \perp as output without passing through a non-null intermediate configuration state. TQ-IH answered negatively. ✓

Status: DERIVED from L-RUN, D7, and R4. No Kolmogorov machinery required. AX-1 derivability pathway now open.

Remark R-TQ — External Confirmation of TQ-IH (April 2026)

TQ-IH was independently reviewed by an external PhD mathematician.

Confirmation: "As long as you have $c_1 \neq c_0$, this implies $c_i \neq c_0$ for some $1 \leq i < n$, at least for $n \geq 2$. This has nothing to do with computation or anything at all, it's a triviality."

This confirms two things. First, TQ-IH holds. Second, and more significantly, the result is domain-independent: it does not depend on any Turing-specific machinery, any Kolmogorov complexity argument, or any property of the semilattice. It holds in any setting where a non-null first step $c_1 \neq c_0$ is given. The derivation from L-RUN alone is complete and unassailable.

Status: EXTERNALLY CONFIRMED — April 2026.

Candidate Theorem T-BUF — Incompressibility Forces Non-Null Execution State

Statement: At the incompressibility threshold P_0 , the Binary Snap $\perp \rightarrow \epsilon_0$ is a structural consequence of execution, not an external trigger.

Step 1 — P_0 identifies the configuration x at which $K(x|n)/n = 1$: the configuration string is incompressible. (D1)

Step 2 — An incompressible configuration at P_0 is informationally extreme (L-INF): its surprisal is unbounded — no finite external program bounds its informational content. A configuration with unbounded informational content has no finite external interpreter and cannot be a static description. Therefore the configuration at P_0 is a live machine state. The design principle connecting informational extremity to forced execution is DA-1 (ZP-E § I-DA1, citing L-INF).

Candidate Theorem T-BUF — Incompressibility Forces Non-Null Execution State

Step 3 — Any execution passes through c_1 (L-RUN). $c_1 \neq \perp$ (L-RUN conclusion).

Step 4 — In (L, \vee, \perp) , this non-null configuration state is $c_1 = \perp \vee \epsilon_0$. By ZP-A D2, this is the Binary Snap.

Conclusion: At P_0 , execution is structurally guaranteed. Execution guarantees a non-null configuration state. That state is ϵ_0 in the semilattice. The derivation pathway is open: $P_0 + \text{L-RUN} + \text{TQ-IH} + \text{ZP-A D2}$, pending cross-framework integration via DA-1 in ZP-E. ✓

Status: CANDIDATE THEOREM — structurally complete within ZP-C. The step from informational extremity (L-INF) to forced execution (DA-1) is a design principle, not a mathematical consequence. Full derivation owned by ZP-E.

Remark R5 — Updated Status of AX-1

Prior to v1.4: AX-1 (Binary Snap Causality) was labeled Axiomatic in ZP-C.

As of v1.4: AX-1 is a Candidate Theorem. The derivation pathway: P_0 (D1) identifies the threshold. L-RUN establishes that execution at the threshold constitutes a non-null state change. TQ-IH establishes that no program avoids this. ZP-A D2 establishes that a non-null state change from \perp is the Binary Snap.

Remaining work: DA-1 (Definitional Alignment) must formally tie instantiation of P_0 to an execution event. This is owned by ZP-E.

Status label: CANDIDATE THEOREM — gap identified and named (DA-1). Closed in ZP-E DA-1 insert.

VI. Open Items Register

Item	Status	Description
S1: Distribution stipulation	Closed — T1	T1 derives P and Q from AX-B1 and RP-1.
OQ-C1: Non-conservation	Closed — T2 rebuilt	Telescoping critique resolved; infinite divergence proven within extended D6.
Smooth embedding, MO-1, P1	Retired	Remain retired from v1.2. Inconsistent with ZP-B.
RP-1: Representation principle	Principle — explicit	Bridge between AX-B1 and probabilistic tools.
RP-2: Branching measure on $Q_2 \setminus \{0\}$	Principle — explicit	Canonical branching measure; representational commitment; required by T2 and L-INF.
CC-2: $c_0 = \perp$	Conditional Claim	Modeling commitment — c_0 identified with \perp in semilattice. Parallel to CC-1 in ZP-A. Required by L-RUN Step 4.
D7: Machine configuration	Defined	Foundation for L-RUN and TQ-IH.
L-RUN: Hardware Lemma	Derived — Lemma	Execution is a non-null state change. Derived from AX-B1 and D7.

Item	Status	Description
TQ-IH: Test Question	Closed — Confirmed	No program can output \perp without a non-null intermediate configuration state. Proven by L-RUN. Externally confirmed April 2026 (R-TQ): domain-independent, requires no Turing-specific or Kolmogorov machinery.
T-BUF: Buffer Overflow Theorem	Candidate Theorem	Incompressibility forces non-null execution state. DA-1 bridge in ZP-E closes this fully.
AX-1: Binary Snap Causality	Candidate Theorem	Derivation pathway formalized. Closed as T-SNAP in ZP-E DA-1 insert.

VII. Validation Status

Component	Status / Notes
$K(x n)$ and P_0 (D1)	Valid — standard algorithmic IT
R1: Scope of P_0	Valid — updated: AX-1 now Candidate Theorem, not bare axiom
RP-1: Representation Principle	Valid — Principle; explicit bridge; resolves reviewer gap in T1
T1: Distributions from AX-B1 + RP-1	Valid — Derived; reviewer gap closed
D2: JSD definition	Valid — standard
T1b: $E = \text{JSD} = 1$ bit	Valid — Derived from AX-B1 + RP-1
D3: Dirac measure δ_0	Valid — standard; compatible with discrete Q_2 topology
R3: Smooth embedding retired	Valid — Retired; inconsistent with ZP-B
D4: Discrete surprisal $I(x)$	Valid — pointwise on $Q_2 \setminus \{0\}$; branching measure stated
RP-2: Branching measure	Valid — Principle; explicit representational commitment; canonical binary branching measure
CC-2: $c_0 = \perp$	Valid — Conditional Claim; modeling commitment; parallel to CC-1 in ZP-A; load-bearing for L-RUN Step 4
D5: Difference operator DF	Valid — antisymmetric; no smoothness assumed
D6: Circulation (extended)	Valid — finite and infinite cases both defined; finite conservation acknowledged
T2: Non-conservation rebuilt	Valid — Derived; telescoping critique addressed; divergence at $\sigma \rightarrow 0$ established
R2: Hamming cross-validation	Valid — Consistency check
D7: Machine configuration	Valid — Defined; standard Turing model; no physics imported
L-RUN: Hardware Lemma	Valid — Derived from AX-B1 and D7

Component	Status / Notes
R4: Configuration vs. output independence	Valid — load-bearing distinction for TQ-IH and T-BUF
TQ-IH: Test question answered	Valid — Derived by L-RUN; no Kolmogorov machinery required
T-BUF: Candidate Theorem	Candidate — structurally complete in ZP-C; DA-1 bridge in ZP-E closes fully
R5: AX-1 status updated	Valid — AX-1 is Candidate Theorem; prior Axiomatic status corrected
R-BRIDGE: K vs. 2-adic surprisal	Valid — Remark; states relationship explicitly: distinct measures converging at P_0 ; independence of L-INF and K paths preserved; cross-reference to ZP-E R-AFA for structural convergence argument