THE ZERO PARADOX

ZP-Q: The Frame-Change

$\perp \to \varepsilon_0$ as a change of point of view

Initial: July 2026 | Current: 1.0, July 2026

Synthesis layer, the sequel to ZP-P. The abstract frame-flip schema is proved sorry-free and choice-free in Lean 4 (ForkFrameChange.lean); the per-domain instances are witnessed in their home layers (the valuation instance — the Riemann sphere — in SnapFrameChange.lean / RiemannSphere.lean; the categorical instance in SeamFrameChange.lean). The universal cross-category identity is not a theorem — it meets a proven wall (Cantor; Lawvere.lean). Adds no axiom; the bottom itself is never positively represented.

The Zero Paradox proves the snap $\perp \to \varepsilon_0$ as a theorem; ZP-Q names what kind of act the snap is. Across every layer the bottom shows the same two features: it is a fixed point — a zero — that carries an infinitude ($v_2(0) = \infty$, unbounded surprisal, the self-referential closure), and every map to it runs one way and reverses only under an inversion that swaps the two poles. These are not two facts. They are two faces of one self-dual object, and the inversion is the frame-change. Read from one frame the frame-change is a bridge (it joins the poles continuously); read from the other it is a fence (a boundary the map cannot cross). ZP-Q records the abstract frame-flip schema, its per-domain instances, and the universal form together with the wall that bounds it.

Like ZP-P, this is a synthesis layer, organised in three tiers held to three standards of proof. Tier 1 (Section I) is the abstract schema, proved in Lean over a complete lattice: it extends ZP-P's fork with its second face, the order-duality that swaps the two closures. Tier 2 (Section II) is the instance catalogue: each domain's frame-change is a theorem in its home layer, and the valuation instance is the framework's originating figure, the p-adic Riemann sphere. Tier 3 (Section III) is the universal form and its walls: the order-theoretic universal holds, the categorical (Lawvere) universal is a proven wall, and the cross-domain identity is a type boundary. The fences in Section III are load-bearing — and they turn out to be built from the same diagonal as the bridges.

A note on what is drawn

ZP-Q characterises the bottom apophatically, and the register is itself frame-relative. Seen from outside — from any frame that looks at $\perp$ — there is no interior: nothing sits behind the description waiting to be opened. The boundary is the whole of it. $\perp$ is the seam the frame-change joins, or the tear it leaves, a boundary of measure zero, all edge and no inside (read flat, this is the Quine atom $\perp = \{\perp\}$ from without — a single node, not a region — and the tear on the sphere — a discontinuity, not a region).

Seen from $\perp$'s own frame — inside the self-reference — the same point is the opposite: infinite depth. $v_2(0) = \infty$; the endless regress $\perp \in \perp \in \perp$; unbounded surprisal. No-interior and infinite-interior are the $0 \leftrightarrow \infty$ poles of one coin, swapped by the frame-change: what has no thickness from without has nothing but depth from within. So the bottom is not an object behind the boundary — it is the boundary, whose own thickness is the frame-flip.

Section I: The Frame-Flip Schema (Tier 1)

ZP-P recorded the fork's first face: over a complete lattice a monotone self-map f has a least fixed point (lfp f, the well-founded / inductive closure, μ) and a greatest fixed point (gfp f, the non-well-founded / coinductive closure, ν), and the fork collapses to a single point — the diagonal fixed point $\perp$ — exactly when f has a unique fixed point (fork_collapse_iff). ZP-Q adds the second face. The two closures are the two charts; the frame-change is the order-duality that carries one to the other. It is the abstract form of the concrete frame-changes below: the p-adic inversion $z \mapsto 1/z$, the categorical op-duality, lattice complementation.

Definition: the frame-change (ForkFrameChange.lean)

For a complete lattice α and a monotone self-map $f : \alpha \to \alpha$:

the two charts are lfp f (μ, the well-founded closure) and gfp f (ν, the non-well-founded closure);

the frame-change is the order-duality $f \mapsto$ f.dual on α^{op} — the abstract analogue of $z \mapsto 1/z$, op, and complementation.

The schema states that the frame-change swaps the two charts, and that they coincide exactly at the diagonal fixed point.

Proposition: lfp_dual_eq_gfp / gfp_dual_eq_lfp

lfp (f.dual) = gfp f and gfp (f.dual) = lfp f.

The frame-change (order-duality) carries the μ-closure to the ν-closure and back: the two charts are swapped, not merged. In Lean both are definitional (rfl) — the swap is built into how gfp is defined (sSup of post-fixed points).

Lean purity: [propext, Quot.sound] — choice-free. ✓

Theorem: fork_is_frameflip — the schema spine

Over any complete lattice, bundling both faces: (i) lfp (f.dual) = gfp f and (ii) gfp (f.dual) = lfp f (the frame-change swaps the charts), together with (iii) lfp f = gfp f $\leftrightarrow \exists! x, f\,x = x$ (the fork collapses at the diagonal fixed point — fork_collapse_iff).

This is the domain-independent form of "the snap is the change of frame": the two charts are swapped by the frame-change, and the bottom is where they meet.

Lean purity: [propext, Quot.sound] — choice-free. ✓

The frame-flip schema is choice-free. fork_is_frameflip and both duality lemmas depend only on [propext, Quot.sound] — propositional extensionality and quotient soundness, the benign kernel axioms used throughout Mathlib. None depends on the Axiom of Choice. As in ZP-P, the conceptual skeleton needs no choice; choice enters only in the analytic realisations (Section II). See AxiomProfile.lean.

Section II: The Instance Catalogue (Tier 2)

The same frame-flip appears across the framework, each domain realising it with its own concrete involution. ZP-Q does not re-prove each instance — each is a theorem in its home layer — it records the catalogue and pins the two instances with a full frame-flip (valuation, category) to their Lean witnesses. The two poles of each row are the two charts; the frame-change is the involution that swaps them.

Domain	The two poles (charts)	Frame-change — witness
Valuation	$0 = \bot \leftrightarrow \infty$	$z \mapsto 1/z$ (rInv) — snap_is_frameflip
Category	initial (μ) ↔ terminal (ν)	op-duality — catseam_is_frameflip
Set theory	Foundation (μ) ↔ AFA (ν)	the fork — fork_collapse_iff
Computability	code ↔ self-application	the diagonal — AbstractSelfApp.unique_fp
State / Hilbert	$X \leftrightarrow X \boxplus X$	the biproduct diagonal — biprod_diagonal_only_zero
Information	measure 0 ↔ surprisal ∞	$-\log$ — info_bottom_diverges
Reals	— (no discrete floor)	inversion breaks at 0 — NO-GO (ZP-F)

The valuation instance is the originating figure of the framework: the p-adic Riemann sphere. Under the tower-rank encoding (P8.lean) the ω-tower climbing to ε_0 lands on the 2-adic floor 0 — so in that chart ε_0 is realised as the bottom. Viewed through the inversion rInv (the one-point compactification of $\mathbb{Q}_2$, RiemannSphere.lean), the same tower rises to ∞ — in that chart ε_0 is realised as the ceiling. rInv is the homeomorphism that swaps the poles $0 \leftrightarrow \infty$.

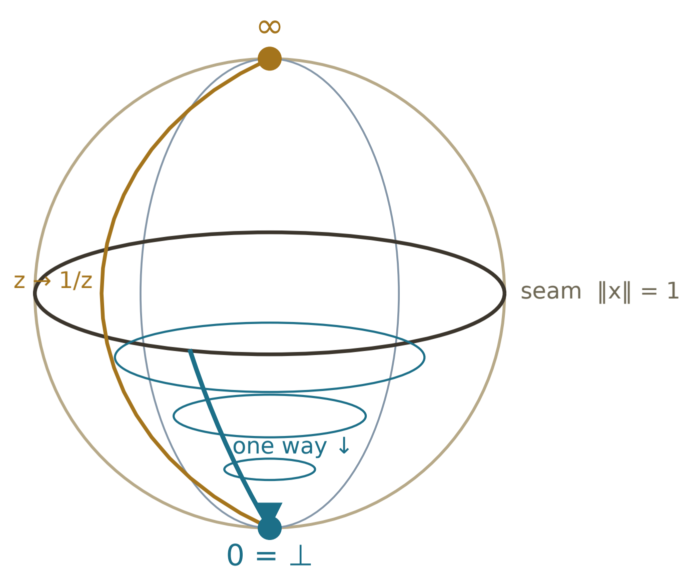

Figure 1. The p-adic Riemann sphere. The poles $0 = \bot$ (the floor) and ∞ (the ceiling) are antipodal; the equator $\|x\| = 1$ is the seam. Inside $\mathbb{Z}_p$ (the lower hemisphere) the ω-tower falls one way to 0; the inversion $z \to 1/z$ (rInv) swaps the poles $0 \leftrightarrow \infty$ — continuous on the sphere, torn at 0 in the field.

Theorem: snap_is_frameflip (SnapFrameChange.lean)

The one ω-tower has two limits, one per chart: it falls to the floor $0 = \perp$ in the encoding chart (ε_0 as $\perp$, cnf_encode_tower_tendsto_zero), and rises to the antipode ∞ in the rInv chart (ε_0 as the ceiling, snap_frameflip_tower_tendsto_infty), with rInv the frame-change swapping the poles (rInv_swaps: $0 \leftrightarrow \infty$).

The descent to $\perp$ and the ascent to the ceiling are the same tower under the frame-change — the valuation-frame realisation of the snap as a change of frame.

Lean purity: [propext, Classical.choice, Quot.sound] — choice-carrying (Classical.choice from the Riemann-sphere continuity machinery). ✓

Remark: continuity is frame-relative — the bridge and the fence

RiemannSphere.lean proves the content that makes the sphere the right figure. In the affine field $\mathbb{Q}_2$ the inversion $z \mapsto 1/z$ is discontinuous at 0 (1/0 is undefined): the map tears. Adjoining ∞ and sending $0 \mapsto \infty$ repairs the tear — on the sphere the inversion is a continuous homeomorphism (rInvHomeo). So the same point 0 is a discontinuity (a fence) from the field frame and a continuous join to ∞ (a bridge) from the sphere frame. Continuity itself is frame-relative; one chart tears where the other heals. The tear at 0 is the apophatic bottom — the place the map fails, never positively drawn.

Theorem: catseam_is_frameflip (SeamFrameChange.lean)

The zero-object seam Z = fD_functor.obj 0 is the μ-bottom (initial) in one chart and the ν-top (terminal) in the other; the op-duality frame-change swaps the two — carrying Z's initial-ness to a terminal-ness of Z^{op} and its terminal-ness to an initial-ness of Z^{op} (terminalOpOfInitial / initialOpOfTerminal). The seam is the op-self-dual zero object.

Off the zero object the coincidence fails (SeamLimColim.lean: generic_object_empty_lim_ne_colim) — as with 0 on the sphere, the frame-change fixes only the bottom.

Lean purity: [propext, Classical.choice, Quot.sound] — choice-carrying (inherited from the ModuleCat / category-theory library). ✓

The remaining instances are referenced, not re-proved here. Set theory: the Quine atom $\perp = \{\perp\}$ is the collapse of the Foundation / AFA fork (ZP-J, ZP-P; fork_collapse_iff, selfMem_eq_singleton_bot). Computability: the bottom is the unique fixed point of self-application (ZP-K; AbstractSelfApp.unique_fp). State / Hilbert: $\perp$ is the unique finite-dimensional fixed point of the biproduct-diagonal $X \mapsto X \boxplus X$ (ZP-H; biprod_diagonal_only_zero). Information: measure runs to 0 at the floor exactly as surprisal runs to ∞ (ZP-C; info_bottom_diverges, addVal_bot), the two joined by $-\log$.

Remark: the reals are the counterexample (the discriminator)

The real numbers are where the frame-change fails, and that is exactly what makes this a diagnostic rather than a pattern read into every domain. $\mathbb{R}$ has no discrete floor to fall to (density), and its inversion is discontinuous at 0 with no repair. Both features of the bottom — the infinitude at the zero and the one-way arrow that reverses under inversion — are absent, and the snap fails (ZP-F). The frame-flip is present precisely where the snap lives and absent precisely where it does not.

Section III: The Universal Form and Its Walls (Tier 3)

Is there a single universal frame-flip — one theorem all the instances are cases of? The question resolves into two halves, and both are proved. At the order-theoretic level the universal holds: fork_is_frameflip (Section I) is a single choice-free theorem over any complete lattice that every instance realises. At the categorical level the universal is a wall. The fences below are not hedging; they mark a genuine boundary, and per-instance frame-flips remain full theorems on either side of it.

The wall: the categorical universal is category-relative (Lawvere.lean)

Lawvere's fixed-point theorem derives a fixed point from a point-surjection $\alpha \to (\alpha \to \beta)$. ZP asks whether each bottom is such a Lawvere-produced fixed point. The answer is category-relative, and it is proved. In Set (all endofunctions) no nontrivial total type carries a Lawvere witness — Cantor forbids it (nontrivial_lattice_no_witness, q2_no_witness) — so the lattice and 2-adic bottoms share the diagonal shape but are provably NOT literal Lawvere instances. Only in the effective category is there a genuine one, where the fixed-point-free diagonal is not computable (computability_face_fixedPoint, Kleene's recursion theorem). So the keystone unifies a shape, not a single mechanism.

Hard fence (permanent): cross-domain identity is a type boundary

The claim that the contact points of all the instances — the p-adic 0, the Quine atom, the Kleene quine, the categorical zero object — are numerically one object is not a theorem and cannot become one. Identity requires a shared type; these are terms of different types in different categories, so "x = y" across them is not false but not well-formed. This is proved at the measure level too: BottomMeasure.lean shows "+∞ at every floor" is a per-domain bundle in incompatible types, not one invariant. The interconnect is therefore point-of-view-specific. What is claimed formally is only that each domain forks, and flips, at its own contact point.

The fence is the bridge. The frame-change that joins the two poles (the bridge) and the diagonal that forbids collapsing them into one object (the fence) are the same diagonal. The bridge only exists between poles that stay two: were they merged there would be nothing to cross. So the boundary is not the price of the connection — it is its condition. The keystone is bounded by the exact mechanism it is built from: the diagonal argument makes the fixed point and, in the same gesture, forbids its universal identification. Per-instance frame-flips are theorems and are not hedged; the fences scope only the universal cross-category claim.

Theorem Summary

Result	Lean source (full path)	Axioms
lfp_dual_eq_gfp	ZeroParadox/Settheory/ForkFrameChange.lean	choice-free
gfp_dual_eq_lfp	ZeroParadox/Settheory/ForkFrameChange.lean	choice-free
fork_is_frameflip	ZeroParadox/Settheory/ForkFrameChange.lean	choice-free
fork_collapse_iff	ZeroParadox/Settheory/FixedPointFork.lean	choice-free
snap_is_frameflip	ZeroParadox/Multihomed/SnapFrameChange.lean	Classical.choice
rInv_swaps	ZeroParadox/Valuation/RiemannSphere.lean	Classical.choice

Result	Lean source (full path)	Axioms
catseam_is_frameflip	ZeroParadox/Category/SeamFrameChange.lean	Classical.choice
nontrivial_lattice_no_witness	ZeroParadox/Category/Lawvere.lean	Classical.choice
computability_face_fixedPoint	ZeroParadox/Category/Lawvere.lean	Classical.choice

Axiom Purity
Frame-flip schema (ForkFrameChange.lean): [propext, Quot.sound] — choice-free. The universal needs no Axiom of Choice; the duality-swap is definitional.
Valuation instance (SnapFrameChange.lean / RiemannSphere.lean): [propext, Classical.choice, Quot.sound] — Classical.choice inherited from the p-adic / one-point-compactification analysis.
Categorical instance (SeamFrameChange.lean): [propext, Classical.choice, Quot.sound] — inherited from the ModuleCat / category-theory library.
The wall (Lawvere.lean): the concrete Cantor no-witness results (nontrivial_lattice_no_witness, q2_no_witness) and computability_face_fixedPoint (Kleene's recursion theorem) carry Classical.choice; the abstract witness machinery (fixedPoint_of_witness, no_witness_of_fixedPointFree) is axiom-free.
Core choice-free, realisation choice-carrying — the framework's standing pattern. Zero sorry. Verified: lake build, July 2026.

End of ZP-Q | The Frame-Change | fork_is_frameflip (choice-free order-theoretic universal) | the Riemann-sphere valuation instance (snap_is_frameflip) and the op-duality categorical instance (catseam_is_frameflip) | the categorical universal is a proven wall (Cantor; Lawvere.lean) | the fence is the bridge | the bottom is never positively represented.