HINGE 6 PREPRINT · V1.0 · APRIL 2026

# The $\tau$-Topos and Its Four-Valued Internal Logic

*Resolving paraconsistent semantic circularity via split-complex idempotents and $\omega$-germ stabilization*

Thorsten Fuchs • Anna-Sophie Fuchs

Correspondence: thorsten@panta-rhei.site

April 2026

DOI: 10.5281/zenodo.19820194

**ABSTRACT**

Building on the earned categorical machine of Hinge 5 [17] and the split-complex boundary algebra $\mathbb{D}$ of Hinge 4 [16], we construct the *$\tau$-topos* $\mathbf{Cat}_\tau$: a countable, boundary- addressed categorical universe equipped with a canonical *four- valued internal logic* $\mathsf{Truth}_4 = B_\sigma(\mathbb{D}) = \{0, e_+, e_-, 1\}$ internalised as the subobject classifier $\Omega_\tau$. We prove: (i) a *topos structure theorem* showing that $\mathbf{Cat}_\tau$ carries all Grothendieck–topos structure (finite limits, exponentials, subobject classifier) earned from $\mathsf{HolEnd}_\tau$ via pre-Yoneda collapse, with $\mathsf{Truth}_4$ as its four truth values $\{\text{Neither, True, False, Both}\} = \{0, e_+, e_-, 1\}$; (ii) a *paraconsistent soundness theorem* showing that the internal logic of $\mathbf{Cat}_\tau$ satisfies the paraconsistent inference rules of Belnap–Priest [2, 27], resolving the classical explosion principle *ex contradictione quodlibet* by the split-complex idempotent structure; (iii) a *circularity- resolution theorem* showing that the classical paradox of self- referential semantics (the Liar, Curry, and the Kleene–Rosser diagonal) admits a constructive stabilisation in $\mathbf{Cat}_\tau$ via $\omega$-germ stabilisation of tails, with *Neither* interpreted ontically as "Cauchy sequences not yet tail-stabilised" rather than epistemically as "unknown truth value"; (iv) a *both-equals-one theorem* showing that Both $= 1$ in $B_\sigma(\mathbb{D})$ is the idempotent unit of $\mathbb{D}$, realising the Hegelian–dialectical "unity of opposites" as the algebraic identity of the Boolean sublattice; and (v) a *hinge-integration theorem* positioning $\mathbf{Cat}_\tau$ in the Panta Rhei bundle and forwarding to Hinge 7's canonical-address NF confluence. The construction is deliberately *countable* and *addressable*: $\mathbf{Cat}_\tau$ has no inaccessible cardinals, no uncountable limits, and no function-space primitives; all categorical content is earned from $\omega$-germ transformers on the countable profinite boundary ring. Lean formalisation is planned in `TauLib.BookII.Topos`.

**Keywords** $\tau$-topos, four-valued internal logic, Truth4, paraconsistent logic, subobject classifier, split-complex idempotent, boundary algebra, $\omega$-germ stabilization, semantic circularity, Neither as ontic uncertainty, Both as idempotent unit, Panta Rhei hinge paper

**MSC** 2020 Mathematics Subject Classification: 18B25, 03B53, 03F65, 18F10, 03G30, 18A40

## CONTENTS

## 1. INTRODUCTION AND STATEMENT OF RESULTS

### 1.1 Position in the Panta Rhei hinge-paper bundle

This paper is **Hinge 6** of the eight-paper Panta Rhei foundational bundle accompanying the 2nd Edition of the series [8, 9, 10]. The bundle consists of seven technical hinges (H1–H7) plus a foundational-anchor paper (H8); in the recommended reading order they are:

**Hinge 1:** *Hyperfactorization* [6] — unique tower-atom decomposition.

**Hinge 2:** *Prime Polarity* [15] — Legendre $(2/p)$ mod-8 prime split.

**Hinge 3:** *Master Constant* $\iota_\tau$ [7] — $\iota_\tau = 2/(\pi + e) \approx 0.341304$.

**Hinge 4:** *The Split-Complex Boundary Algebra* [16] — $\mathbb{D} = \mathcal{R}'_\partial[j]/(j^2 - 1)$ and the four-atom generator dictionary.

**Hinge 5:** *$\tau$-Holomorphy on the Boundary Algebra* [17] — $\omega$-germ transformers, wave-equation Cauchy–Riemann, earned categorical machine, and $\mathsf{HolEnd}_\tau$ via pre-Yoneda collapse.

**Hinge 6:** *The $\tau$-Topos and Its Four-Valued Internal Logic (this paper)* — builds $\mathbf{Cat}_\tau$ with subobject classifier $\Omega_\tau = B_\sigma(\mathbb{D})$ and four-valued paraconsistent internal logic, resolving semantic circularity via $\omega$-germ stabilisation.

**Hinge 7:** *Address Resolution, Not Calculation* [5] — the genealogical DAG, the Cayley word metric, and the ontic ultrametric; supplies the NF confluence theorem that discharges Hinge 5's pre-Yoneda collapse and this paper's topos-structure "modulo Hinge 7" caveats.

**Hinge 8:** *The $\tau$-Kernel as Foundational Architecture* [18] — foundational-anchor paper (also readable as an entry point): ontic identity invariance, diagonal–linear correspondence, $*$-autonomous placement; names what the seven technical hinges collectively earn.

The present paper imports from Hinge 5 the earned categorical machine (composition, identity, associativity, functoriality as theorems), the pre-Yoneda collapse, and $\mathsf{HolEnd}_\tau$; from Hinge 4 the split-complex boundary algebra $\mathbb{D}$ and its Boolean sublattice $B_\sigma(\mathbb{D}) = \{0, e_+, e_-, 1\}$ of $\sigma$-equivariant idempotents. It supplies to Hinge 7 the topos-theoretic frame in which canonical addressability is realised as a genuinely universal (not merely choice-theoretic) normal form.

### 1.2 The paraconsistent semantic circularity problem

Classical topos theory builds its internal logic over a Heyting algebra with two truth values ($\{\bot, \top\}$) or occasionally over a generalised truth-object with finitely many intermediate values. The classical problem of *semantic circularity* — the Liar paradox, Curry's paradox, Kleene–Rosser's diagonal, and the Russell–Tarski hierarchy — is typically resolved by either (a) restricting self-reference syntactically (Tarski's hierarchy of meta-languages), or (b) adopting a three- or four-valued paraconsistent logic in which the principle *ex contradictione quodlibet* fails, so that local contradictions do not trivialise the whole system [2, 27].

Both classical resolutions have a common structural weakness: they are *externally imposed*. The hierarchy is an extra syntactic layer; the four truth values are axiomatic. Neither approach explains *why* those particular resolutions are the right ones in the first place. The present paper offers a third approach:

> *The four truth values of a paraconsistent logic are not external axioms, but the four elements of the $\sigma$-equivariant Boolean sublattice $B_\sigma(\mathbb{D})$ of the split-complex boundary algebra. Semantic circularity is resolved not by restricting self-reference, but by recognising that self-referential fixed points live in the $\omega$-germ stabilisation of a Cauchy tail-sequence, whose stabilised value is a **specific element** of $B_\sigma(\mathbb{D})$.*

Concretely, the four truth values of $\mathsf{Truth}_4$ are identified as

$$\mathsf{Truth}_4 \;:=\; B_\sigma(\mathbb{D}) \;=\; \{0, e_+, e_-, 1\} \;=\; \{\text{Neither, True, False, Both}\},$$

with:

- Neither $= 0 \in \mathbb{D}$: the zero element; semantically, "no tail stabilisation yet". This is *ontic* uncertainty (the Cauchy sequence has not yet settled on its stabilised value), not *epistemic* uncertainty (we do not know its value).
- True $= e_+ \in \mathbb{D}$: the plus-lobe idempotent; semantically, the proposition $p$ holds and $\neg p$ does not.
- False $= e_- \in \mathbb{D}$: the minus-lobe idempotent; semantically, $\neg p$ holds and $p$ does not.
- Both $= 1 \in \mathbb{D}$: the idempotent unit $e_+ + e_- = 1$; semantically, the proposition $p$ and its negation $\neg p$ jointly hold in the $\omega$-germ stabilisation. This is the paraconsistent fixed point of self-referential propositions like "this sentence is false."

The key identification Both $= 1$ — the idempotent unit of $\mathbb{D}$ — realises the classical Hegelian "unity of opposites" as a genuine algebraic identity. The Liar proposition $L \leftrightarrow \neg L$ has its $\omega$-germ fixed point at $[\![L]\!] = 1 = e_+ + e_- \in B_\sigma(\mathbb{D})$, not at the non- existent classical truth value.

### 1.3 Main results

We establish five main theorems. Each is anchored at **[$\tau$-Effective]** (the current paper's primary scope tier); a scope-promotion path to **[Established]** is mapped in §9.2 contingent on Hinge 7's NF confluence and Book II's classical-topos embedding.

**Theorem 1.1** (Topos structure of $\mathbf{Cat}_\tau$ **[$\tau$-Effective]**, modulo Hinge 7 canonical-address NF confluence)**.** *The countable boundary-addressed category $\mathbf{Cat}_\tau$, constructed from the earned categorical machine* $\mathsf{Hol}_\tau$ *and the pre-Yoneda collapse of Hinge 5, carries:*

- *all finite limits (terminal object $1$, binary products in the restricted sense of Definition 3.7, equalisers);*
- *exponential objects $Y^X$ for every pair of admissible carriers $X, Y$, realised as the representable boundary-addressed object of* $\mathsf{Hol}_\tau(X, Y)$*;*
- *a subobject classifier $\Omega_\tau \cong B_\sigma(\mathbb{D}) = \{0, e_+, e_-, 1\}$ with characteristic morphism $\chi_\tau\colon \mathrm{Sub}_\tau \to \Omega_\tau$.*

*In particular $\mathbf{Cat}_\tau$ is an elementary topos in the sense of Lawvere–Tierney, with the distinguishing feature that its subobject classifier is the four-element Boolean sublattice of the split- complex boundary algebra $\mathbb{D}$ rather than the two-element Boolean lattice $\{0, 1\}$ of classical set theory.*

**Theorem 1.2** (Paraconsistent soundness of $\mathsf{Truth}_4$ **[$\tau$-Effective]**)**.** *The internal logic of $\mathbf{Cat}_\tau$, interpreted via the subobject-classifier semantics $[\![-]\!]\colon \mathrm{Prop}_\tau \to \Omega_\tau = B_\sigma(\mathbb{D})$, is sound with respect to the four-valued paraconsistent Belnap–Dunn logic [2]: all classical Boolean-algebraic identities valid on $B_\sigma(\mathbb{D})$ are derivable in the internal logic, while the principle* ex contradictione quodlibet *($p \wedge \neg p \vdash q$) fails because the idempotent unit* Both $= 1$ *is a fixed point of conjunction-with- negation. Consequently $\mathbf{Cat}_\tau$'s internal logic is a genuine paraconsistent logic, not merely a classical logic with extra labels.*

**Theorem 1.3** (Circularity resolution via $\omega$-germ stabilisation **[$\tau$-Effective]**)**.** *Every self-referential proposition $p = \Phi(p)$ arising from a decidable tail-coherent propositional template $\Phi\colon \Omega_{\mathrm{tail}} \to \Omega_{\mathrm{tail}}$ admits a unique $\omega$-germ stabilised truth value $[\![p]\!] \in B_\sigma(\mathbb{D})$, computed as the $\sim$-tail-equivalence class of the Cauchy iteration $\Phi^n(\bot) \in \Omega_{\mathrm{tail}}$. In particular:*

- *The Liar $L = \neg L$ stabilises at $[\![L]\!] =$* Both $= 1$.
- *Curry's paradox $C = (C \to \bot)$ stabilises at $[\![C]\!] =$* Neither $= 0$ *(the iteration fails to stabilise within any finite-witness budget).*
- *A genuine paradox-free fixed point $T = T$ stabilises at $[\![T]\!] \in \{\mathrm{True}, \mathrm{False}\} = \{e_+, e_-\}$ according to its truth-sector membership.*

*The resolution is* constructive*: the classification of $[\![p]\!]$ into one of $\{\mathrm{Neither}, \mathrm{True}, \mathrm{False}, \mathrm{Both}\}$ is a finite-witness decidable procedure on $\Omega_{\mathrm{tail}}$.*

**Theorem 1.4** (Both $= 1$: idempotent unit of the Boolean sublattice **[$\tau$-Effective]**)**.** *Inside $B_\sigma(\mathbb{D}) \subset \mathbb{D}$, the paraconsistent truth value* Both $\in \mathsf{Truth}_4$ *is precisely the multiplicative unit $1$ of the ring $\mathbb{D}$ and the join $e_+ \vee e_-$ of the two lobe idempotents in the Boolean sublattice. Equivalently, the Hegelian "unity of opposites" $p \wedge \neg p$ at a paraconsistent fixed point lands on the idempotent identity of the split-complex algebra, not on a failed truth value.*

### 1.4 Lean roadmap (preview)

Full formalisation is targeted at `TauLib.BookII.Topos` [19] in the Lean 4 proof assistant [25], comprising the following planned modules:

- `TauTopos.lean` — the $\tau$-topos $\mathbf{Cat}_\tau$ and its finite-limit structure.
- `Truth4.lean` — the four-valued internal logic $\mathsf{Truth}_4 = B_\sigma(\mathbb{D})$ and the paraconsistent connectives.
- `SubobjectClassifier.lean` — $\Omega_\tau$ as subobject classifier with characteristic morphism $\chi_\tau$.
- `ExponentialObjects.lean` — $Y^X$ via pre-Yoneda collapse (imports Hinge 5 / Hinge 7 dependencies).
- `Circularity.lean` — constructive fixed-point stabilisation for self-referential propositions.

See §9 for the full proof-chain sketch.

### 1.5 Hinge-integration theorem

**Theorem 1.5** (Hinge 6 integration tabulation **[$\tau$-Effective]**). *Hinge 6 is positioned in the Panta Rhei hinge bundle by the following four backward dependencies and two forward obligations.*

- ***From Hinge 4*** *[16]: the unique boundary algebra $\mathbb{D}$, its four idempotent atoms $\{0, e_+, e_-, 1\}$, the $\sigma$-involution, and the Boolean sublattice structure of $B_\sigma(\mathbb{D})$.*
- ***From Hinge 5*** *[17]: the earned categorical machine $\mathsf{Hol}_\tau$, the pre-Yoneda collapse $y_\tau\colon \tau \to \mathbf{PSh}_\tau$, and the holomorphic endomorphism category $\mathsf{HolEnd}_\tau$ — the categorical host on which $\mathbf{Cat}_\tau$ is built.*
- ***Backward to Hinges 1–3****: the ABCD coordinates [6], the prime-polarity $B/C$ split [15], and the master constant $\iota_\tau$ [7] realise as specific propositional transformers in $\mathrm{Prop}_\tau$ or as $\sigma$-fixed global sections of $\Omega_\tau$.*
- ***To Hinge 7*** *(forthcoming): the topos-theoretic frame provided here becomes the categorical universe in which canonical addressability is realised as the universal property of a specific terminal-addressable object in $\mathbf{Cat}_\tau$. The NF confluence theorem of Hinge 7 is the technical lemma on which $\mathbf{Cat}_\tau$'s exponential-object construction depends.*

*The present paper completes the bundle's* sixth *hinge: the internal-logic layer in which the $\sigma$-equivariant idempotent sublattice $B_\sigma(\mathbb{D})$ is realised as the subobject classifier of a genuine elementary topos.*

## 2. PRELIMINARIES: IMPORTS FROM HINGES 4 AND 5

### 2.1 Boundary algebra $\mathbb{D}$ and its $\sigma$-equivariant Boolean sublattice

We recall, without re-proving, the split-complex boundary algebra of Hinge 4 [16]. Throughout, $\mathcal{R}'_\partial = \mathcal{R}_\partial[1/2]$ denotes the dyadic localisation of the countable profinite boundary ring $\mathcal{R}_\partial \cong \varprojlim_k \mathbb{Z}/M_k\mathbb{Z}$ over the primorial ladder $(M_k)$ [16, §§3–4].

**Remark 2.1** (Imported: split-complex boundary algebra **[Established]**). The *split-complex boundary algebra* is the free commutative $\mathcal{R}'_\partial$-algebra on one generator $j$ modulo $j^2 = +1$:

$$\mathbb{D} \;:=\; \mathcal{R}'_\partial[j] \,/\, (j^2 - 1), \tag{1}$$

free of rank 2 as an $\mathcal{R}'_\partial$-module with basis $\{1, j\}$, so every $z \in \mathbb{D}$ is uniquely $z = a + j\,b$ with $a, b \in \mathcal{R}'_\partial$. For background on the classical split-complex ("hyperbolic") number plane see [32, 29, 4]; the $\tau$-framework's $\mathbb{D}$ is the boundary-addressable, dyadically localised counterpart of those classical algebras. The elliptic alternative $\mathbb{C}$ arising from $i^2 = -1$ is the classical several-complex-variables setting [20, 28, 21, 24]; the $\tau$-topos of this paper is by construction the hyperbolic counterpart, with paraconsistent internal logic replacing the Heyting logic of the classical holomorphy topos. Hinge 4 Theorem 1.6 [16] proves that $\mathbb{D}$ is the unique commutative $\mathcal{R}'_\partial$-algebra satisfying four $\tau$-kernel structural constraints (binary rank, commutativity, two nontrivial orthogonal idempotents, canonical polarity-swap involution); the elliptic alternative $\mathcal{R}'_\partial[i]/(i^2+1)$ is excluded by [16, Thm. 1.7].

**Remark 2.2** (Imported: idempotent basis **[Established]**). The canonical orthogonal idempotents

$$e_+ \;:=\; \tfrac{1}{2}(1 + j), \qquad e_- \;:=\; \tfrac{1}{2}(1 - j), \tag{2}$$

satisfy $e_+ + e_- = 1$, $e_+ \cdot e_- = 0$, $e_+^2 = e_+$, $e_-^2 = e_-$, $j = e_+ - e_-$, and induce an internal direct-sum splitting

$$\mathbb{D} \;\cong\; \mathcal{R}'_\partial \cdot e_+ \,\oplus\, \mathcal{R}'_\partial \cdot e_-, \qquad z \;=\; z_+\,e_+ + z_-\,e_-, \tag{3}$$

with lobe-projection $\mathcal{R}'_\partial$-algebra homomorphisms $\pi_\pm\colon \mathbb{D} \to \mathcal{R}'_\partial$ [16, §7].

**Remark 2.3** (Imported: $\sigma$-involution and Boolean sublattice **[Established]**)**.** $\mathbb{D}$ carries a canonical involutive $\mathscr{R}'_\partial$-algebra automorphism $\sigma\colon \mathbb{D} \to \mathbb{D}$ with

$$\sigma(j) = -j,\ \sigma(e_+) = e_-,\ \sigma(e_-) = e_+,\ \sigma(1) = 1,\ \sigma(0) = 0,\ \sigma^2 = \mathrm{id}_{\mathbb{D}}\,. \tag{4}$$

The *$\sigma$-equivariant Boolean sublattice*

$$B_\sigma(\mathbb{D}) \ :=\ \{0, e_+, e_-, 1\} \ \subset\ \mathbb{D}, \tag{5}$$

the smallest Boolean sub-$*$-algebra of $\mathrm{Idem}(\mathbb{D})$ closed under $\sigma$ and containing $\{e_+, e_-\}$, has exactly four elements [16, Thm. 1.8]. These realise the four-atom dictionary $0 \leftrightarrow \alpha$-null, $e_+ \leftrightarrow \gamma$ (EM), $e_- \leftrightarrow \eta$ (strong), $1 \leftrightarrow \alpha$-total (gravity), with the non-idempotent $\sigma$-fixed crossing scalar $\iota_\tau = 2/(\pi + e)$ as the $\omega$-mediator outside $B_\sigma(\mathbb{D})$.

**Remark 2.4** (Imported: $\sigma$-fixed subalgebra is the real axis **[Established]**)**.** Writing $z = z_+ e_+ + z_- e_-$, the condition (4) gives $\sigma(z) = z_- e_+ + z_+ e_-$, so $\sigma(z) = z$ iff $z_+ = z_-$:

$$\mathbb{D}^\sigma \ :=\ \{z \in \mathbb{D} : \sigma(z) = z\} \ =\ \mathscr{R}'_\partial \cdot 1 \ =\ \mathscr{R}'_\partial \cdot (e_+ + e_-) \tag{6}$$

is the *real axis* of $\mathbb{D}$ [17, Cor. 8.5]. The real axis $\mathbb{D}^\sigma$ hosts $\iota_\tau$ and will host the paraconsistent truth value Both $= 1$ in §5.

**Remark 2.5** (Truth-sector labelling)**.** For use in §§4–7 we fix the bijective labelling

$$\text{Neither} \ :=\ 0, \qquad \text{True} \ :=\ e_+, \qquad \text{False} \ :=\ e_-, \qquad \text{Both} \ :=\ 1. \tag{7}$$

At this stage (7) is a naming convention; the paraconsistent semantics that earns these labels is developed in §§5–7, and the identification Both $= 1$ as the idempotent unit of $\mathbb{D}$ (Theorem 1.4) is proved there.

### 2.2 Earned categorical machine from Hinge 5

We recall, without re-proving, the earned categorical machine of Hinge 5 [17]. All objects below are countable and finite-witness decidable [17, Rem. 2.2, 3.6].

**Remark 2.6** (Imported: $\omega$-tails and tail-equivalence **[Established]**)**.** The type $\Omega_{\text{tail}}$ of *$\omega$-tails* consists of infinite coherent prefix chains over the $\tau$-native token alphabet [17, Def. 2.1]. For each depth $k \in \mathrm{Idx}$ there is a decidable prefix-agreement relation $\equiv_k$ on $\Omega_{\text{tail}}$; the *full tail-equivalence*

$$t \sim t' \quad \Longleftrightarrow \quad \forall k \in \mathrm{Idx},\ t \equiv_k t' \tag{8}$$

is reflexive, symmetric, and transitive [17, Lem. 2.3]. Each $X \in \mathrm{Obj}(\tau)$ comes with a $\sim$-invariant decidable predicate $\mathsf{Tail}_X\colon \Omega_{\text{tail}} \to \mathsf{Prop}$.

**Remark 2.7** (Imported: admissibility predicates **[Established]**)**.** A *tail transformer code* $c \in \mathsf{Code}$ has intensional semantics $[\![c]\!]\colon \Omega_{\text{tail}} \to \Omega_{\text{tail}}$ given by primitive-recursive rewriting [17, Def. 3.1, Rem. 3.7]. The three admissibility predicates on $c$ relative to carriers $X, Y$ are:

$$\mathsf{Typed}(X, Y, c) \ :\Longleftrightarrow\ \ \forall t,\ \mathsf{Tail}_X(t) \Rightarrow \mathsf{Tail}_Y([\![c]\!](t)); \tag{9}$$

$$\mathsf{Stable}(X, Y, c) \ :\Longleftrightarrow\ \ \mathsf{Tail}_X(t) \wedge \mathsf{Tail}_X(t') \wedge t \sim t' \ \Rightarrow\ [\![c]\!](t) \sim [\![c]\!](t'); \tag{10}$$

$$\text{tail-indep.} \ :\Longleftrightarrow\ \ \exists k_0,\ \mathsf{Tail}_X(t) \wedge \mathsf{Tail}_X(t') \wedge t \equiv_{k_0} t' \ \Rightarrow\ [\![c]\!](t) \sim [\![c]\!](t'). \tag{11}$$

The type of *$\tau$-holomorphic maps* from $X$ to $Y$ is

$$\mathsf{Hol}_\tau(X, Y) \ :=\ \{c \in \mathsf{Code} \mid \mathsf{Typed}(X, Y, c) \wedge \mathsf{Stable}(X, Y, c) \wedge \text{tail-indep.}\}, \tag{12}$$

modulo $\sim$-equivalence of codes. No Cartesian product $X \times Y$, function graph, or set-theoretic function space is used at this layer [17, Rem. 3.9]; countability of $\mathsf{Hol}_\tau(X, Y)$ is immediate from finite-witness decidability.

**Remark 2.8** (Imported: earned categorical machine **[Established]**, modulo Hinge 7 for strong confluence). From the admissibility predicates (9)–(11) alone, [17, Thm. 1.6, §7] derives as theorems:

(a) *Earned composition.* If $c \in \mathsf{Hol}_\tau(X, Y)$ and $c' \in \mathsf{Hol}_\tau(Y, Z)$, then the sequential action $[\![c']\!] \circ [\![c]\!]$ is the semantics of an admissible code $c' \circ c \in \mathsf{Hol}_\tau(X, Z)$.

(b) *Earned identity.* Each carrier $X$ has a canonical tail-fixing NF code $\mathrm{id}_X \in \mathsf{Hol}_\tau(X, X)$ that is the two-sided unit of composition.

(c) *Earned associativity.* Composition is associative up to NF equivalence of codes; weak confluence holds at this layer and is strengthened to strong confluence by Hinge 7 (§2.3).

(d) *Functoriality.* $X \mapsto \mathsf{Hol}_\tau(X, -)$ is a functor on the probe category $P_\tau$ of primorial-depth refinements.

No category axioms are imposed; all four clauses are theorems.

**Remark 2.9** (Imported: pre-Yoneda collapse **[τ-Effective]**, modulo Hinge 7). The *probe category* $P_\tau$ has as objects primorial-depth carrier types $X_n$ and as morphisms refinement transformers $X_n \to X_m$ ($n \geq m$); it is thin and posetal. Under the $\tau$-native admissibility predicates, the Yoneda functor

$$y_\tau \colon \tau \longrightarrow \mathbf{PSh}_\tau := [\, P_\tau^{\mathrm{op}}, \mathbf{Set}\,], \qquad X \mapsto \mathsf{Hol}_\tau(-, X) \tag{13}$$

*collapses at the presheaf level*: every $\mathsf{Hol}_\tau(-, X)$ is representable by a canonical boundary-addressed code $c_X \in \partial\tau^3$, where $\partial\tau^3 := \partial(\tau^1 \times_f T^2)$ is the boundary of the Hinge 1 fibered product [6, 17]. The assignment $X \mapsto c_X$ is well-defined up to canonical NF equivalence and functorial in $P_\tau$-refinements [17, Thm. 9.7]; the strict form (strict uniqueness of $c_X$) is the forward dependency on Hinge 7 stated in §2.3.

**Remark 2.10** (Imported: holomorphic endomorphism category $\mathsf{HolEnd}_\tau$ **[τ-Effective]**). The *holomorphic endomorphism category*

$$\mathsf{HolEnd}_\tau \;:=\; \{(X, f) : X \in \mathrm{Obj}(\tau),\; f \in \mathsf{Hol}_\tau(X, X)\} \tag{14}$$

has as morphisms $\phi \colon (X, f) \to (Y, g)$ the admissible intertwiners $\phi \in \mathsf{Hol}_\tau(X, Y)$ with $g \circ \phi = \phi \circ f$ in NF [17, Def. 9.1, Prop. 9.2]. Identity is $\mathrm{id}_X$; composition is earned; associativity and unit laws are inherited from Remark 2.8. Under the pre-Yoneda collapse, $\mathsf{HolEnd}_\tau$ is concretely representable as the category of boundary-addressed pairs $(c_X, f_X) \in \partial\tau^3$ [17, Prop. 9.10].

**Remark 2.11** (Imported: $\sigma$-equivariant refinement $\mathsf{HolEnd}_\tau^\sigma$ **[τ-Effective]**). The $\sigma$-involution (4) lifts to $\mathsf{Hol}_\tau(X, X)$ by $\bar{f} := \sigma_X \circ f \circ \sigma_X$, which is again admissible [17, Thm. 8.2]. The *$\sigma$-equivariant endomorphism monoid*

$$\mathsf{HolEnd}_\tau^\sigma \;:=\; \{(X, f) \in \mathsf{HolEnd}_\tau : \bar{f} = f\} \tag{15}$$

(with morphisms inherited from $\mathsf{HolEnd}_\tau$, restricted to $\sigma$-equivariant intertwiners) is closed under identity and composition, hence a wide subcategory of $\mathsf{HolEnd}_\tau$ [17, Prop. 9.14]; it is the natural host for $\sigma$-fixed structure, in particular for $\iota_\tau$ as the universal $\sigma$-fixed scalar [17, Thm. 9.17].

**Remark 2.12** (Imported: idempotent-supported holomorphy **[τ-Effective]**). Every admissible map into $\mathbb{D}$ factors through the two lemniscate lobes: there is an internal direct-sum decomposition of $\mathsf{Hol}_\tau(X, \mathcal{R}'_\partial)$-modules

$$\mathsf{Hol}_\tau(X, \mathbb{D}) \;=\; e_+ \cdot \mathsf{Hol}_\tau(X, \mathcal{R}'_\partial) \;\oplus\; e_- \cdot \mathsf{Hol}_\tau(X, \mathcal{R}'_\partial), \tag{16}$$

with a unique pair $(f_+, f_-) \in \mathsf{Hol}_\tau(X, \mathcal{R}'_\partial)^2$ for each $f \in \mathsf{Hol}_\tau(X, \mathbb{D})$ satisfying $f = e_+ f_+ + e_- f_-$ [17, Thm. 8.5]; the sector components are the lobe projections $f_\pm = \pi_\pm \circ f$. This transformer-level splitting drives the classification of self-referential stabilised fixed points into the four truth sectors (7) in §7.

### 2.3 Forward dependency on Hinge 7 (canonical-address NF confluence)

Hinges 5 and 6 both depend on a normalisation theorem whose full proof lies in Hinge 7 [17, 16]. For self-containment, we state it here as an assumed black-box lemma.

**Lemma 2.13** (Canonical-address NF confluence **[$\tau$-Effective]**, assumed modulo Hinge 7)**.** *For every admissible* $c \in \mathsf{Code}$ *with* $\mathsf{Typed}(X, Y, c) \wedge \mathsf{Stable}(X, Y, c) \wedge$ *tail-indep., the normalisation operator* $\mathsf{Norm}\colon \mathsf{Code} \to \mathsf{Code}$ *defined by iterated application of the primitive-recursive NF rewrite rules of [17, §7.3] terminates in finitely many steps and produces a* unique *canonical normal form* $\mathsf{Norm}(c) \in \mathsf{Code}$ *independent of the rewriting path. Equivalently, the NF reduction system on admissible codes is strongly normalising and confluent.*

**Remark 2.14** (Status of Lemma 2.13 and scope caveat)**.** Weak confluence on already-admissible codes is proven in [17, §7.2]; this suffices for pre-Yoneda representability up to NF equivalence (Remark 2.9). The strong (strict uniqueness) form of Lemma 2.13 is the subject of Hinge 7 (*Address Resolution, Not Calculation*), established using the genealogical DAG and the Cayley word metric [6]. Statements of the present paper that ultimately depend on strict NF uniqueness — specifically the exponential-object and subobject-classifier clauses of Theorem 1.1 and the NF normalisation of Both to 1 in Theorem 1.4 — are flagged **[$\tau$-Effective]** *modulo Hinge 7*. The black-box character of Lemma 2.13 here is deliberate scope-tier discipline, not a gap.

**Remark 2.15** (What Lemma 2.13 buys)**.** With Lemma 2.13 granted, the pre-Yoneda collapse becomes a *strict* representability theorem: the assignment $X \mapsto c_X$ picks out a strictly unique address in $\partial\tau^3$, and the exponential object $Y^X$ of §3 becomes a functor into the category of strict boundary-addressed codes rather than a functor into NF-equivalence classes. Hinge 6 is written to work in either regime: the topos structure theorem holds up to NF equivalence unconditionally, and upgrades to a strict version once Hinge 7 discharges Lemma 2.13.

### 2.4 Notation conventions specific to this paper

We fix the notation used in §§3–7 for the $\tau$-topos, its subobject classifier, and its internal logic. All symbols below are local to Hinge 6; they expand into objects already earned in Hinges 4–5 or into objects to be constructed in the body of this paper.

**Definition 2.16** (The $\tau$-topos $\mathbf{Cat}_\tau$ **[$\tau$-Effective]**, modulo Hinge 7)**.** *The* $\tau$-topos $\mathbf{Cat}_\tau$ *is the countable boundary-addressed category with objects the admissible carriers* $X \in \mathrm{Obj}(\tau)$ *and morphisms the* $\tau$*-holomorphic maps* $\mathsf{Hol}_\tau(X, Y)$*, endowed with the finite-limit structure, exponential objects, and subobject classifier constructed in §§3–4. The full construction is deferred to §3; the notation* $\mathbf{Cat}_\tau$ *is fixed here for forward reference.*

**Definition 2.17** ($\Omega_\tau$, $\mathbf{PSh}_\tau$, $\mathbf{Sh}_\tau$ **[$\tau$-Effective]**)**.** *Write* $\Omega_\tau$ *for the subobject classifier of* $\mathbf{Cat}_\tau$ *(identified in §4 as* $\Omega_\tau \cong B_\sigma(\mathbb{D})$ *via Remark 2.3) and* $\mathbf{PSh}_\tau$ *for the presheaf category* $[\,P_\tau^{\mathrm{op}}, \mathbf{Set}\,]$ *of Remark 2.9. The notation* $\mathbf{Sh}_\tau$ *is reserved for the sheaf subcategory to be taken up in Book II [9]; it does not appear in the present paper's constructive core.*

**Definition 2.18** (Four-valued internal logic $\mathsf{Truth}_4$ **[$\tau$-Effective]**)**.**

$$\mathsf{Truth}_4 \;:=\; B_\sigma(\mathbb{D}) \;=\; \{0, e_+, e_-, 1\} \;=\; \{\text{Neither, True, False, Both}\}, \tag{17}$$

*with the truth-sector labelling of* (7)*. At this preliminary stage* $\mathsf{Truth}_4$ *is a set of four symbols; the paraconsistent Belnap–Dunn lattice structure (meet, join, entailment, negation) on* $\mathsf{Truth}_4$ *is defined in §5 and proven sound in §6.*

**Definition 2.19** ($\mathrm{Sub}_\tau$, $\chi_\tau$, $[\![-]\!]$ **[$\tau$-Effective]**)**.** *For each* $X \in \mathrm{Obj}(\mathbf{Cat}_\tau)$, $\mathrm{Sub}_\tau(X)$ *denotes the poset of subobjects of* $X$ *(monomorphisms modulo iso), and* $\chi_\tau\colon \mathrm{Sub}_\tau(-) \to \Omega_\tau$ *the natural transformation classifying subobjects via the universal property of* $\Omega_\tau$ *(§4). A subobject* $m\colon S \hookrightarrow X$ *is classified by its characteristic morphism* $\chi_\tau(m)\colon X \to \Omega_\tau$ *with values in* $B_\sigma(\mathbb{D})$*, interpreted via* (7)*. For a proposition* $p$ *in the internal language of* $\mathbf{Cat}_\tau$*, the* propositional-interpretation bracket

$$[\![p]\!] \;\in\; \Omega_\tau = B_\sigma(\mathbb{D}) \tag{18}$$

*is the truth value of* $p$ *in the subobject-classifier semantics; the interpretation rules for* $\wedge$, $\vee$, $\vdash$*, and paraconsistent implication are the subject of §§5–6.*

**Remark 2.20** (Entailment and inference-rule notation)**.** The symbol $\vdash$ denotes internal entailment on $\Omega_\tau$; the macro $\dfrac{\text{premises}}{\text{conclusion}}$ is used in §6 to display the paraconsistent Belnap–Dunn rule schemata [2, 27]. On $B_\sigma(\mathbb{D})$, meet $\wedge$ and join $\vee$ coincide with the orthogonal-idempotent calculus of (2): $e_+ \wedge e_- = 0$ (Boolean-lattice $\bot$) and $e_+ \vee e_- = 1$ (Boolean-lattice $\top$), with $0$ and $1$ as the $\bot$ and $\top$ of $(B_\sigma(\mathbb{D}), \wedge, \vee)$.

**Remark 2.21** (Scope-tier discipline for this paper). Imports in §§2.1–2.2 inherit their scope tags from their source papers [16, 17]; new definitions local to this paper (§2.4) are **[$\tau$-Effective]**; pure classical-algebra corollaries drawn from $B_\sigma(\mathbb{D})$ viewed as an abstract Boolean sublattice are **[Established]**; and any theorem requiring strict canonical-address NF confluence is flagged **[$\tau$-Effective]** *modulo Hinge 7* per Remark 2.14. The five main theorems of §1 (Theorems 1.1–1.5) all lie at **[$\tau$-Effective]** under this discipline; promotion to **[Established]** awaits Hinge 7 and the classical-topos embedding tabulated in Book II [9].

## 3. THE $\tau$-TOPOS $\mathbf{Cat}_\tau$: OBJECTS, MORPHISMS, AND FINITE LIMITS

### 3.1 Definition of $\mathbf{Cat}_\tau$

We now assemble the category that will carry the topos structure of this paper. The assembly is not a fresh axiomatisation: every piece of categorical data has already been *earned* as theorem in Hinge 5 [17], and every addressable piece of structure has been canonicalised through the pre-Yoneda collapse of [17, §9]. The role of the present section is to package those results into a single category $\mathbf{Cat}_\tau$ and to verify that it carries the finite-limit structure required of an elementary topos.

Two guiding constraints shape the definition below. First, the *diagonal discipline* [17, §6] forbids us from taking arbitrary Cartesian products of carriers: a pair of admissible carriers does not, in general, yield an admissible product carrier. Consequently the notion of binary product in $\mathbf{Cat}_\tau$ is *typed* (Definition 3.7) rather than free. Second, the absence of any pre-existing ambient category means that $\mathbf{Cat}_\tau$ must be *concretely addressable*: every object and every morphism must pin down a canonical code in the boundary algebra $\partial\tau^3$ of Hinge 1 [6]. The pre-Yoneda collapse of Hinge 5 [17, Thm. 9.7] supplies exactly that address.

**Definition 3.1** (The $\tau$-topos $\mathbf{Cat}_\tau$). *II.D.TOPOS.CATTAU The* $\tau$-topos $\mathbf{Cat}_\tau$ *is the category whose data are:*

- ***Objects.*** *Admissible carriers* $X \in \mathrm{Obj}(\tau)$ *in the sense of Hinge 5 ([17, Def. 3.1]): for each candidate carrier $X$, the tail-predicate* $\mathsf{Tail}_X$ *is certified (decidable prefix predicates* $\mathsf{Pref}_{k,\sigma}$ *witnessing membership at every primorial depth), and $X$ is equipped with the canonical* boundary address $c_X \in \partial\tau^3$ *supplied by the pre-Yoneda collapse ([17, Thm. 9.7]).*
- ***Morphisms.*** *For each ordered pair of objects* $X, Y \in \mathrm{Obj}(\mathbf{Cat}_\tau)$, *the morphism set is*

$$\mathbf{Cat}_\tau(X,Y) \;:=\; \mathsf{Hol}_\tau(X,Y) \;=\; \{\,[c] \;:\; \mathsf{Typed}(X,Y,c) \wedge \mathsf{Stable}(X,Y,c) \wedge \textit{tail-indep}(c)\,\}/\!\sim,$$

  *i.e. $\sim$-classes of NF-coded tail transformers satisfying the three admissibility predicates of [17, §3]. Each class contains a canonical NF representative* $c_f \in \mathsf{Code}$ *and, via pre-Yoneda collapse, a canonical boundary address in* $\partial\tau^3$.
- ***Composition.*** *The* earned composition *of [17, Thm. 7.2]: for* $f \in \mathbf{Cat}_\tau(X,Y)$ *and* $g \in \mathbf{Cat}_\tau(Y,Z)$,

$$g \circ f \;:=\; [\mathsf{Norm}(c_g \cdot c_f)] \;\in\; \mathbf{Cat}_\tau(X,Z).$$

- ***Identity.*** *The* earned identity *of [17, Def. 7.4, Prop. 7.5]: for each $X$,* $\mathrm{id}_X := [c_{\mathrm{id}_X}]$ *is the $\sim$-class of the empty NF code.*

*Associativity and two-sided unit laws hold by [17, Thm. 7.8, Prop. 7.6], so* $\mathbf{Cat}_\tau$ *is a category in the usual sense.*

**Remark 3.2** (What $\mathbf{Cat}_\tau$ inherits from $\mathbf{Hol}_\tau$). The full subcategory of $\mathbf{Cat}_\tau$ spanned by the admissible carriers of Hinge 5 is exactly the earned category $\mathbf{Hol}_\tau$ of [17, Cor. 7.14]. The passage from $\mathbf{Hol}_\tau$ to $\mathbf{Cat}_\tau$ is thus a *labelling* step, not a fresh construction: we read the earned category as an elementary categorical universe and will — in §§4–5 — equip it with the subobject classifier and four-valued internal logic that distinguishes it from $\mathbf{Hol}_\tau$. No new morphism type is introduced in the present section.

**Remark 3.3** (Three structural hallmarks). Three features set $\mathbf{Cat}_\tau$ apart from classical categorical universes and must be kept in view throughout the remainder of this paper.

- (H1) **Countability.** $\mathbf{Cat}_\tau$ is a countable category (Theorem 3.19 below). There are no inaccessible cardinals, no uncountable limits, and no hidden set-theoretic universe. The bound is structural: it comes from the countable profinite boundary ring ([16, Thm. 1.1]).
- (H2) **Boundary-addressedness.** Every object and every morphism in $\mathbf{Cat}_\tau$ lifts to a canonical address in $\partial\tau^3$ (Theorem 3.20). There is no "object up to isomorphism" whose representative cannot be written down as an NF code on the boundary algebra.

(H3) **No function-space primitives.** Exponential objects $Y^X$ are *not* axiomatically included. Where they exist (§4–§7 reconstruct them via pre-Yoneda collapse), their existence is an earned theorem that cites Hinge 5's categorical machine and Hinge 7's forthcoming NF confluence. The present section is explicit about function spaces being *unavailable as primitives*: the diagonal discipline of [17, Thm. 6.5] is still in force.

### 3.2 Terminal object

Before addressing binary products, we verify that $\mathbf{Cat}_\tau$ has a *terminal object*. Terminality is the lightest-weight finite limit: it requires only the existence of a carrier and of a unique morphism into it from every object. Both hold straightforwardly because the $\omega$-tail formalism of [17, §3] already admits a one-point carrier whose tail predicate is trivially certified.

**Definition 3.4** (The boundary-addressable singleton). *Let $1 := \{*\}_\tau$ denote the carrier whose only admissible tail is the constant $\omega$-tail at the normal-form zero element $0 \in \mathbb{D}$; equivalently, the trivial $\sim$-class of NF codes representing the constant-tail code with a single output symbol. We call $1$ the* boundary-addressable singleton.

The admissibility of $1$ is immediate: $\mathsf{Tail}_1(t)$ holds iff $t$ is $\sim$-equal to the constant zero tail, which is a decidable prefix predicate of depth $0$. The pre-Yoneda address of $1$ is the degenerate boundary address $c_1 = 0 \in \partial\tau^3$, which is the unique idempotent minimum of the boundary algebra.

**Theorem 3.5** ($1$ is terminal in $\mathbf{Cat}_\tau$ [$\tau$-Effective]). *II.T.TOPOS.TERM For every $X \in \mathrm{Obj}(\mathbf{Cat}_\tau)$ there exists a unique morphism $!_X : X \to 1$ in $\mathbf{Cat}_\tau$. Consequently $1$ is a terminal object of $\mathbf{Cat}_\tau$.*

*Lean-grade sketch. Existence.* For each admissible carrier $X$, let $c_{!_X}$ be the constant-tail NF code that sends every $\omega$-tail satisfying $\mathsf{Tail}_X$ to the constant zero tail. The three admissibility predicates are satisfied:

- $\mathsf{Typed}(X, 1, c_{!_X})$: for $t$ with $\mathsf{Tail}_X(t)$, the image $[\![c_{!_X}]\!](t)$ is the constant zero tail, which by construction satisfies $\mathsf{Tail}_1$.
- $\mathsf{Stable}(X, 1, c_{!_X})$: if $t \sim t'$ with $\mathsf{Tail}_X(t)$, $\mathsf{Tail}_X(t')$, then both images equal the constant zero tail, which is trivially $\sim$-equivalent to itself.
- Tail-independence beyond depth $0$: the output does not depend on any prefix of the input.

Thus $c_{!_X} \in \mathsf{Hol}_\tau(X, 1)$ is an admissible transformer, and its $\sim$-class is a morphism $!_X : X \to 1$ in $\mathbf{Cat}_\tau$.

*Uniqueness.* Let $g \in \mathbf{Cat}_\tau(X, 1)$ be any morphism. By $\mathsf{Typed}(X, 1, c_g)$, the semantics $[\![c_g]\!] : \Omega_{\text{tail}} \to \Omega_{\text{tail}}$ sends every $t$ with $\mathsf{Tail}_X(t)$ to a tail satisfying $\mathsf{Tail}_1$. But $\mathsf{Tail}_1$ admits only the single $\sim$-class of the constant zero tail. Hence for every admissible $t$ the output $[\![c_g]\!](t)$ is $\sim$-equal to $[\![c_{!_X}]\!](t)$. Because $\mathsf{Hol}_\tau$-morphisms are $\sim$-classes, $g = !_X$ in $\mathbf{Cat}_\tau(X, 1)$. Uniqueness follows.

Existence plus uniqueness is the universal property of a terminal object. □

**Remark 3.6** (Terminality is cheap; universality is not). The terminal object in $\mathbf{Cat}_\tau$ is structurally trivial because the singleton tail-predicate has exactly one $\sim$-class. The interesting content of the definition is not the existence of $1$ but the observation that $1$'s boundary address is canonically $0 \in \partial\tau^3$: this will be used in §4 to anchor the "False" pole of the subobject classifier $\Omega_\tau$.

### 3.3 Typed binary products

We now turn to binary products. Here the diagonal discipline of Hinge 5 [17, §6] asserts its structural grip. An arbitrary pair $X, Y$ of admissible carriers does *not* yield an admissible Cartesian product carrier $X \times Y$: the product-and-slice construction presupposes a classical set-theoretic stratum that the $\tau$-kernel declines to admit (clause (DD1) of [17, Def. 6.2]). Forcing it would trigger the integral-domain collapse traced in [17, Thm. 6.5] and would destroy the split-complex idempotent pair $(e_+, e_-)$ on which the whole framework stands.

What *is* available is a notion of *typed product*: a pair carrier that exists for pairs which are admissibility-compatible in a specific, decidable sense. The construction is forced by the pre-Yoneda collapse [17, §9.4]: when two admissible carriers factor jointly through the boundary address $\partial\tau^3$, their pair address — the concatenation of their boundary codes — is itself admissible, and yields a product carrier. When they do not, no pair carrier exists, and any morphism trying to land on such a pair is structurally inadmissible.

**Definition 3.7** (Typed binary product $X \times^\tau Y$). *II.D.TOPOS.TYPEDPROD Let $X, Y \in \mathrm{Obj}(\mathbf{Cat}_\tau)$. A* typed binary product *of X and Y is a triple $(P, \pi_1, \pi_2)$ consisting of an object $P \in \mathrm{Obj}(\mathbf{Cat}_\tau)$ and two morphisms $\pi_1 : P \to X$, $\pi_2 : P \to Y$ satisfying:*

(TP1) ***Pair code.*** *There exists an NF code $c_\times \in \mathsf{Code}$ encoding the pair carrier P, such that $\mathsf{Typed}(P, X, c_{\pi_1})$ and $\mathsf{Typed}(P, Y, c_{\pi_2})$ hold where $c_{\pi_1}, c_{\pi_2}$ are the canonical first- and second-coordinate projection codes built from $c_\times$.*

(TP2) ***Joint tail-coherence.*** *The boundary addresses $c_X, c_Y \in \partial\tau^3$ are* jointly tail-coherent*: there exists a primorial depth $k_0$ beyond which the two addresses admit a common cylinder refinement that is admissible in $\partial\tau^3$.*

(TP3) ***Universal property.*** *For every $Z \in \mathrm{Obj}(\mathbf{Cat}_\tau)$ and every pair $f : Z \to X$, $g : Z \to Y$ such that the graph-pair $(c_f, c_g)$ is itself* jointly tail-coherent *(i.e. the two boundary addresses of f and g admit a common cylinder refinement), there exists a unique $\langle f, g\rangle : Z \to P$ in $\mathbf{Cat}_\tau$ such that $\pi_1 \circ \langle f, g\rangle = f$ and $\pi_2 \circ \langle f, g\rangle = g$.*

*When it exists, the typed product is written $X \times^\tau Y$ (with $\pi_1, \pi_2$ implicit from context).*

**Remark 3.8** (What "typed" means concretely). The adjective "typed" in *typed binary product* refers to the joint-tail-coherence condition (TP2, TP3), which is a *decidable* predicate on pair addresses: by [16, Thm. 1.1], the countable profinite boundary ring admits a finite-witness test for whether two addresses share a common cylinder refinement at any fixed primorial depth. Hence the notion of typed product is not a vague categorical recipe but a concrete decidable constraint on the admissibility predicates. Pairs that fail joint coherence have *no* product carrier; the category $\mathbf{Cat}_\tau$ is not cartesian but *partially cartesian* in the sense made precise below.

**Remark 3.9** (The diagonal-discipline shadow). Definition 3.7 is the positive face of the diagonal discipline. Under (DD1)–(DD4) of [17, Def. 6.2], the free Cartesian product $X \times Y$ is never an admissible carrier. What Definition 3.7 says is that the *constrained* Cartesian product — the one restricted to boundary-coherent pairs — is available, because the boundary-coherent condition provides exactly the additional structural datum that makes admissibility closure work. The constrained product is the $\tau$-native analogue of the Cartesian product, but it is not free: each typed product requires the joint-coherence certificate.

**Theorem 3.10** (Existence of typed products for admissible pairs **[$\tau$-Effective]**, modulo Hinge 7 canonical-address NF confluence). *II.T.TOPOS.TYPEDPROD-EXISTS Let $X, Y \in \mathrm{Obj}(\mathbf{Cat}_\tau)$ be admissible carriers whose boundary addresses $c_X, c_Y \in \partial\tau^3$ are jointly tail-coherent (condition (TP2) of Definition 3.7). Then the typed binary product $X \times^\tau Y$ exists in $\mathbf{Cat}_\tau$ and is unique up to canonical NF-equivalence.*

*Lean-grade sketch. Construction.* Let $k_0$ be the primorial depth at which $c_X, c_Y$ admit a common cylinder refinement. Define the pair code

$$c_\times \;:=\; \mathsf{Norm}(c_X \uplus c_Y),$$

where $\uplus$ denotes the admissible *address concatenation* in the boundary algebra $\partial\tau^3$, gluing the two address trees along their common $k_0$-cylinder refinement. This is the $\tau$-native "pair operation" on boundary addresses, well-defined by [16, Thm. 1.1] whenever joint coherence holds, and reduced to canonical NF by the rewriting system of Hinge 5 (with confluence discharged forward to Hinge 7's NF confluence theorem). Let $P$ be the admissible carrier coded by $c_\times$; its tail predicate is the joint-cylinder predicate of $c_X$ and $c_Y$ at depth $k_0$.

*Projections.* The first-coordinate projection $\pi_1 : P \to X$ is coded by the canonical *forget-right* NF operation on the pair address: $c_{\pi_1} := \mathsf{Norm}(\pi_1^* \cdot c_\times)$, where $\pi_1^*$ is the elementary rewriting step that erases the $c_Y$-component at each cylinder refinement. Typing, stability, and tail-independence follow from the admissibility of $c_X$ and of the elementary rewriting $\pi_1^*$, using the earned-composition machine of [17, Thm. 7.2] applied to $\pi_1^* \cdot c_\times$. Similarly for $\pi_2 : P \to Y$.

*Universal property.* Let $Z \in \mathrm{Obj}(\mathbf{Cat}_\tau)$ and let $f : Z \to X$, $g : Z \to Y$ be morphisms whose boundary addresses $c_f, c_g$ are jointly tail-coherent. The pair morphism

$$\langle f, g\rangle \;:=\; [\mathsf{Norm}(c_f \uplus c_g)]$$

is an admissible $Z \to P$ morphism by the same admissibility-closure argument. The projection identities $\pi_1 \circ \langle f, g\rangle = f$ and $\pi_2 \circ \langle f, g\rangle = g$ follow from the NF confluence of the forget-and-concatenate elementary operations on boundary addresses:

concatenating followed by forgetting the second coordinate reduces to the identity on the first address, using [17, Thm. 7.8] (associativity) as the backing lemma.

*Uniqueness.* If $h\colon Z \to P$ satisfies $\pi_1 \circ h = f$ and $\pi_2 \circ h = g$, then the boundary address of $h$ agrees with $\mathsf{Norm}(c_f \uplus c_g)$ up to $\sim$, because the pair of coordinate conditions is jointly strong enough to pin down the full pair address up to canonical NF. Hence $h = \langle f, g\rangle$ as $\mathbf{Cat}_\tau$-morphisms. □

**Remark 3.11** (Non-typed products are excluded)**.** The converse of Theorem 3.10 is the diagonal-discipline theorem of [17, Thm. 6.5]: if $c_X, c_Y$ are not jointly tail-coherent, then no admissible pair carrier exists, and in particular the arbitrary Cartesian product $X \times Y$ is not in $\mathbf{Cat}_\tau$. This is not a missing feature but a structural constraint: the discipline is what protects the split-complex idempotent pair $(e_+, e_-)$, on which the four-valued internal logic of §5 depends.

**Corollary 3.12** (Admissible diagonals via typed self-products **[$\tau$-Effective]**)**.** *For every $X \in \mathrm{Obj}(\mathbf{Cat}_\tau)$, the self-pair $(c_X, c_X)$ is trivially jointly tail-coherent (both addresses coincide), so the typed self-product $X \times^\tau X$ exists. The diagonal $\Delta_X\colon X \to X \times^\tau X$ is the typed pairing $\langle \mathrm{id}_X, \mathrm{id}_X\rangle$ of Definition 3.7.*

*Proof.* $(c_X, c_X)$ is trivially jointly tail-coherent at any primorial depth. Apply Theorem 3.10 to obtain $X \times^\tau X$ and Definition 3.7(TP3) to obtain $\Delta_X = \langle \mathrm{id}_X, \mathrm{id}_X\rangle$. □

**Remark 3.13** (What the diagonal *does not* give us)**.** Corollary 3.12 may look like a restoration of the classical diagonal $\Delta\colon X \to X \times X$, but the typed-product qualifier is essential: the diagonal is admissible only at *boundary-coherent self-pairs*, and the target $X \times^\tau X$ is not the free Cartesian square of $X$. In particular, free contraction of a typed hypothesis is still blocked by clause (DD4) of the diagonal discipline ([17, Def. 6.2]): the diagonal exists as a morphism in $\mathbf{Cat}_\tau$, but it cannot be used to duplicate a token of $X$ in the meta-logic.

### 3.4 Equalisers

We turn to equalisers. Given parallel morphisms $f, g\colon X \to Y$ in $\mathbf{Cat}_\tau$, the equaliser is the maximal sub-carrier of $X$ on which $f$ and $g$ agree as $\sim$-classes of NF codes. Equalisers exist unconditionally in $\mathbf{Cat}_\tau$: unlike binary products, they do not require a joint tail-coherence condition, because they are constructed from a *single* carrier $X$ and its admissible sub-carriers.

**Definition 3.14** (Equalising sub-carrier)**.** *Let $f, g\colon X \to Y$ in $\mathbf{Cat}_\tau$. An* equalising sub-carrier *of $(f, g)$ is a carrier $E \subseteq_\tau X$ together with an inclusion $e\colon E \hookrightarrow X$ (an admissible transformer whose code is the canonical inclusion of $\mathsf{Tail}_E$ into $\mathsf{Tail}_X$) such that*

$$f \circ e \;=\; g \circ e \qquad \text{in } \mathbf{Cat}_\tau(E, Y).$$

*The "$\subseteq_\tau$" symbol denotes sub-carrier-at-the-boundary: $\mathsf{Tail}_E$ is a decidable sub-predicate of $\mathsf{Tail}_X$ corresponding to a cylinder refinement of the boundary address $c_X$.*

**Theorem 3.15** (Equalisers exist in $\mathbf{Cat}_\tau$ **[$\tau$-Effective]**)**.** *II.T.TOPOS.EQ For every parallel pair $f, g\colon X \to Y$ in $\mathbf{Cat}_\tau$, there exists an equaliser $e\colon E \hookrightarrow X$ in the sense of Definition 3.14 satisfying the universal property: for every $h\colon Z \to X$ with $f \circ h = g \circ h$, there exists a unique $\tilde{h}\colon Z \to E$ with $e \circ \tilde{h} = h$.*

*Lean-grade sketch.* *Construction of $E$.* Define the tail-predicate

$$\mathsf{Tail}_E(t) \;:\Longleftrightarrow\; \mathsf{Tail}_X(t) \;\wedge\; [\![c_f]\!](t) \sim [\![c_g]\!](t).$$

This is a decidable predicate on $\Omega_{\text{tail}}$: the conjunction is decidable because $\mathsf{Tail}_X$ is decidable (from admissibility of $X$) and $\sim$-equality of output tails is decidable at every finite primorial depth, using the tail-independence bounds of $c_f$ and $c_g$ to reduce the comparison to a finite witness depth $k_0 = \max(k_f, k_g) + \delta$ with $\delta$ the combined tail-shift of $f$ and $g$.

The predicate $\mathsf{Tail}_E$ is a cylinder sub-refinement of $\mathsf{Tail}_X$ at depth $k_0$, so it corresponds to an admissible cylinder sub-address $c_E \in \partial\tau^3$ sitting under $c_X$. This makes $E$ an admissible carrier of $\mathbf{Cat}_\tau$.

The inclusion $e\colon E \hookrightarrow X$ is the canonical forget-constraint NF code: $c_e$ is the empty NF code on $\Omega_{\text{tail}}$-data restricted to $\mathsf{Tail}_E$, so $\mathsf{Typed}(E, X, c_e)$, $\mathsf{Stable}(E, X, c_e)$, and tail-independence at depth 0 hold trivially (the code is the identity on the underlying $\omega$-tails; the only structural content is the type-restriction of its domain). Hence $e \in \mathbf{Cat}_\tau(E, X)$.

*Equalising property.* By construction, $[\![c_f]\!](t) \sim [\![c_g]\!](t)$ for every $t$ with $\mathsf{Tail}_E(t)$. Hence the $\sim$-classes of $f \circ e$ and $g \circ e$ are equal at every admissible input, i.e. $f \circ e = g \circ e$ as $\mathbf{Cat}_\tau$-morphisms.

*Universal property.* Let $h\colon Z \to X$ satisfy $f \circ h = g \circ h$. Then for every $t$ with $\mathsf{Tail}_Z(t)$, the image $[\![c_h]\!](t)$ satisfies $\mathsf{Tail}_X$ and (by hypothesis, reading the equality of $\mathbf{Cat}_\tau$-morphisms at the tail level) $[\![c_f]\!]([\![c_h]\!](t)) \sim [\![c_g]\!]([\![c_h]\!](t))$. Hence $[\![c_h]\!](t)$ satisfies $\mathsf{Tail}_E$. So the NF code $c_h$ factors through $c_E$: define $c_{\tilde{h}} := c_h$ (the same code, but now typed $Z \to E$), and verify the three admissibility predicates at the refined target: typing follows from the pointwise $\mathsf{Tail}_E$-membership of the image; stability and tail-independence are inherited from $c_h$'s admissibility in $\mathsf{Hol}_\tau(Z, X)$. Thus $\tilde{h} := [c_{\tilde{h}}] \in \mathbf{Cat}_\tau(Z, E)$ and $e \circ \tilde{h} = h$ by construction.

*Uniqueness.* If $\tilde{h}'\colon Z \to E$ also satisfies $e \circ \tilde{h}' = h$, then at the tail level $[\![c_{\tilde{h}'}]\!] = [\![c_h]\!]$ (because $e$ acts as the identity on $\Omega_{\text{tail}}$-data), so $c_{\tilde{h}'} \sim c_h = c_{\tilde{h}}$ and $\tilde{h}' = \tilde{h}$ as $\mathbf{Cat}_\tau(Z, E)$-morphisms. □

**Remark 3.16** (Equalisers are cheaper than products)**.** Equalisers do not require joint tail-coherence because they are sub-carrier constructions on a *single* source $X$. The refinement $\mathsf{Tail}_E \subseteq \mathsf{Tail}_X$ is an intersection in the decidable predicate lattice, not a concatenation of boundary addresses, so it is always admissible. This is the formal statement of the widely repeated slogan "in $\mathbf{Cat}_\tau$, restriction is free but pairing is constrained." Equalisers are free; products are typed.

### 3.5 Finite limits

We can now combine terminal object, typed binary products, and equalisers to obtain all finite limits of admissibility-compatible diagrams.

**Theorem 3.17** ($\mathbf{Cat}_\tau$ has all finite limits **[$\tau$-Effective]**, modulo Hinge 7 canonical-address NF confluence)**.** *II.T.TOPOS.FINLIM For every finite diagram $D\colon J \to \mathbf{Cat}_\tau$ whose images are pairwise jointly tail-coherent (in the sense of Definition 3.7(TP2)), the limit* $\lim D \in \mathbf{Cat}_\tau$ *exists and is unique up to canonical NF-equivalence. In particular, $\mathbf{Cat}_\tau$ admits all finite products of admissible pair-tuples, pullbacks of admissible cospans, and equalisers of admissible parallel pairs.*

*Lean-grade sketch.* By a standard categorical reduction (Mac Lane, Cat. Work. Math., Thm. V.2.1), a category admits all finite limits iff it admits a terminal object, binary products, and equalisers. The three ingredients have been constructed:

- Terminal object: Theorem 3.5.
- Typed binary products (at jointly tail-coherent pairs): Theorem 3.10.
- Equalisers (unconditionally): Theorem 3.15.

For finite $n$-ary products, iterate the typed binary product: if the $n$-tuple of carriers $X_1, \dots, X_n$ is jointly tail-coherent (i.e. the $n$ boundary addresses admit a common cylinder refinement at some primorial depth), then each nested typed product $(((X_1 \times^\tau X_2) \times^\tau X_3) \cdots) \times^\tau X_n$ is admissible by repeated application of Theorem 3.10. Unique factorisation through the product is then forced by $n$-ary uniqueness of the pair code.

For pullbacks of an admissible cospan $X \xrightarrow{f} Z \xleftarrow{g} Y$: if $X, Y$ are jointly tail-coherent, form $X \times^\tau Y$ and take the equaliser of the two composites $X \times^\tau Y \rightrightarrows Z$ (namely $f \circ \pi_1$ and $g \circ \pi_2$) via Theorem 3.15. The resulting sub-carrier is the pullback $X \times^\tau_Z Y$.

Finite general limits reduce to finite products and equalisers by the classical construction, which applies verbatim in $\mathbf{Cat}_\tau$ once one restricts to jointly tail-coherent diagram-images. □

**Remark 3.18** (Joint tail-coherence as the typed-cone constraint)**.** The joint-tail-coherence hypothesis in Theorem 3.17 is the structural remnant of the diagonal discipline at the level of finite diagrams. For arbitrary finite diagrams on arbitrary pairs, limits need not exist in $\mathbf{Cat}_\tau$; they exist precisely when the diagram-images admit a boundary-address common refinement. In the forthcoming topos structure of §4, every subobject-classifying diagram is automatically jointly tail-coherent (a consequence of the pre-Yoneda collapse's canonicity), so the typed-cone constraint does not restrict the internal-logic constructions that follow.

### 3.6 Countability and boundary-addressedness

We now make explicit the two global structural properties of $\mathbf{Cat}_\tau$ that distinguish it from classical categorical universes and that will be used repeatedly in §§4–8.

**Theorem 3.19** (Countability of $\mathbf{Cat}_\tau$ **[$\tau$-Effective]**). *II.T.TOPOS.COUNT Both the object class and the morphism class of $\mathbf{Cat}_\tau$ are countable:*

$$|\mathrm{Obj}(\mathbf{Cat}_\tau)| \;\leq\; \aleph_0, \qquad |\,\mathrm{Mor}(\mathbf{Cat}_\tau)| \;\leq\; \aleph_0.$$

*Equivalently, $\mathbf{Cat}_\tau$ is a* small *category in the set-theoretic sense, and fits entirely within the countable ordinal $\omega$.*

*Lean-grade sketch. Objects.* Every object $X \in \mathrm{Obj}(\mathbf{Cat}_\tau)$ is an admissible carrier with canonical boundary address $c_X \in \partial\tau^3$. The boundary algebra $\partial\tau^3$ is a countable profinite algebra by [16, Thm. 1.1], so the set of admissible boundary addresses is countable. Pre-Yoneda collapse [17, Thm. 9.7] makes the assignment $X \mapsto c_X$ injective up to canonical NF-equivalence, so $\mathrm{Obj}(\mathbf{Cat}_\tau)$ embeds injectively into a countable set.

*Morphisms.* Every morphism is a $\sim$-class of NF codes in $\mathsf{Code}$, and $\mathsf{Code}$ is a countable type (finite syntactic structures over a finite alphabet of elementary rewriting steps). The $\sim$-quotient at most reduces the cardinality. Hence for any fixed pair $(X, Y)$, the hom-set $\mathbf{Cat}_\tau(X, Y)$ is countable; and because the set of pairs is countable, the total morphism class is countable by a countable union of countable sets.

*Total cardinality.* $|\,\mathrm{Mor}(\mathbf{Cat}_\tau)| \leq \aleph_0 \cdot \aleph_0 = \aleph_0$, and $|\mathrm{Obj}(\mathbf{Cat}_\tau)| \leq \aleph_0$. □

**Theorem 3.20** (Boundary-addressedness of $\mathbf{Cat}_\tau$ **[$\tau$-Effective]**, modulo Hinge 7 canonical-address NF confluence). *II.T.TOPOS.BDDADR There exist canonical assignments*

$$\mathrm{addr}_{\mathrm{obj}} : \mathrm{Obj}(\mathbf{Cat}_\tau) \;\longrightarrow\; \partial\tau^3, \qquad \mathrm{addr}_{\mathrm{mor}} : \mathrm{Mor}(\mathbf{Cat}_\tau) \;\longrightarrow\; \partial\tau^3,$$

*each injective up to canonical NF equivalence, such that:*

(i) *for each $X \in \mathrm{Obj}(\mathbf{Cat}_\tau)$, $\mathrm{addr}_{\mathrm{obj}}(X) = c_X$ is the canonical boundary code of $X$ from the pre-Yoneda collapse;*

(ii) *for each $f \in \mathbf{Cat}_\tau(X, Y)$, $\mathrm{addr}_{\mathrm{mor}}(f) = c_f$ is the canonical NF code of $f$, which concatenates with $c_X$ and $c_Y$ via the admissible boundary operations $\uplus$ of Theorem 3.10;*

(iii) *composition and identity are computed at the address level:* $\mathrm{addr}_{\mathrm{mor}}(g \circ f) = \mathsf{Norm}(\mathrm{addr}_{\mathrm{mor}}(g) \cdot \mathrm{addr}_{\mathrm{mor}}(f))$ *and* $\mathrm{addr}_{\mathrm{mor}}(\mathrm{id}_X) = c_{\mathrm{id}_X}$ *the empty NF code.*

*In particular, $\mathbf{Cat}_\tau$ is concretely representable in the boundary algebra: every categorical operation reduces to an NF rewriting on boundary addresses.*

*Lean-grade sketch.* (i) is the pre-Yoneda collapse [17, Thm. 9.7]. (ii) is its morphism-level extension [17, Prop. 9.10], which assigns to each $f \in \mathsf{Hol}_\tau(X, Y)$ the address obtained by its NF code. (iii) reduces composition and identity at the address level to the earned composition and identity of [17, Thm. 7.2, Prop. 7.5], evaluated on the boundary side of the pre-Yoneda embedding. Injectivity up to canonical NF equivalence is the uniqueness clause of the pre-Yoneda collapse and the Church–Rosser property of the NF rewriting system (discharged forward to Hinge 7's NF confluence theorem). □

**Remark 3.21** (Structural payoff of boundary-addressedness). Theorem 3.20 is the load-bearing lemma behind the internal-logic constructions of §§4–7. Because every object and every morphism lives at a *finite- witness decidable address*, the subobject classifier of §4 can be evaluated pointwise on boundary addresses (no Grothendieck-descent machinery required), and the circularity-resolution theorem of §7 can compute the stabilised truth value of a self-referential proposition as a *constructive* primorial-depth iteration on NF codes. Without boundary-addressedness, these constructions would require either an ambient set-theoretic universe (which the $\tau$-framework declines to admit) or Grothendieck-topos machinery (which the countability of $\mathbf{Cat}_\tau$ renders unnecessary).

### 3.7 Non-elementary features that $\mathbf{Cat}_\tau$ does not carry

We close the structural-backbone section with an explicit enumeration of what $\mathbf{Cat}_\tau$ does *not* carry as native structure. Each item will be either (a) earned later in this paper (§§4–5) via a construction that cites the machinery already in hand, or (b) forwarded to Book II [9] as beyond the scope of the present hinge paper.

- (N1) **No uncountable limits.** Because $\mathbf{Cat}_\tau$ is countable (Theorem 3.19), it does not carry limits indexed by uncountable diagrams. This is not an oversight: the countability bound is structural (from [16, Thm. 1.1]) and ensures that the subobject classifier $\Omega_\tau$ of §4 is itself a countable object. Uncountable limits — were they permitted — would force the boundary algebra to become uncountable, breaking the countable profinite structure on which the four-atom idempotent dictionary depends.
- (N2) **No function-space primitives.** Exponential objects $Y^X$ are not part of the present structural backbone. Where the topos structure of Theorem 1.1 requires them, they are *earned* via the pre-Yoneda collapse: $Y^X$ is realised as the representable boundary-addressed object whose boundary address is the canonical code of $\mathsf{Hol}_\tau(X, Y)$. The construction is completed in §4 (for the $X = 1$ case), generalised to arbitrary $X$ via the standard $\mathrm{Hom}(-, Y^X) \cong \mathrm{Hom}(- \times^\tau X, Y)$ adjunction in §5, and discharged to Hinge 7's NF confluence for associativity of the induced application map. In the present section we simply note that function spaces are *not* primitive.
- (N3) **No classical two-element subobject classifier $\{0, 1\}$.** The classical topos-theoretic subobject classifier assigns two truth values ($\top, \bot$) to a subobject, forming a two-element Boolean algebra. In $\mathbf{Cat}_\tau$, the *correct* subobject classifier is the four-element Boolean sublattice $B_\sigma(\mathbb{D}) = \{0, e_+, e_-, 1\}$ of the split-complex boundary algebra $\mathbb{D}$ — this is $\Omega_\tau$, developed in §4. The two extra truth values ($e_+$ = True / $e_-$ = False as distinct from the idempotent unit/zero) capture the $\sigma$-equivariant idempotent split forced by the diagonal discipline ([17, Thm. 6.5]). A hypothetical two-valued classifier would collapse the split and destroy the $B/C$-sector decomposition of Hinge 3 [7].
- (N4) **No raw power-object $\mathcal{P}(X)$.** The free power object is unavailable because it would require the free exponential $\Omega_\tau^X$, which by (N2) is not primitive. Power objects are earned in §4 only as the representables of admissible subobject-classifying diagrams.
- (N5) **No free contraction in the meta-logic.** The meta-logical discipline (DD4) of [17, Def. 6.2] is still in force: a typed token cannot be copied silently in the ledger of $\mathbf{Cat}_\tau$-statements. This has a concrete effect on the internal logic (§5): the internal "and" is not idempotent in the free sense because the resource accounting of $B_\sigma(\mathbb{D})$ tracks whether a proposition $p$ has been asserted once or twice. What looks like $p \wedge p$ in the external meta-language unfolds as a typed pair at the internal level, and that pair must be jointly tail-coherent to exist as a subobject.
- (N6) **No ambient set-theoretic universe.** $\mathbf{Cat}_\tau$ is not a small subcategory of some larger classical universe such as **Set** or a Grothendieck topos. It is the universe. Classical comparisons — e.g., the embedding $\mathbf{Cat}_\tau \hookrightarrow \mathbf{Set}^{P_\tau^{\mathrm{op}}}$ provided by the forgetful functor on boundary addresses — exist and are useful (§8), but they do not implicitly import a larger universe into the definition.

**Remark 3.22** (Scope-tier reminder). All statements in this section are at the **[$\tau$-Effective]** tier, modulo the Hinge 7 canonical-address NF confluence dependency flagged in Theorems 3.10, 3.17, and 3.20. The definitions (Definitions 3.1, 3.7, 3.14) are unlabelled: they are definitions, not scope-tagged theorems. Classical-algebra facts cited from Hinges 4 and 5 (countable profiniteness of $\partial\tau^3$; admissibility closure under NF concatenation; pre-Yoneda collapse) are tagged in their source-hinge citations; here we import them as lemmas.

**Remark 3.23** (Forward outlook). With $\mathbf{Cat}_\tau$ constructed, finite limits verified, countability and boundary-addressedness recorded, and non-elementary features flagged, the structural backbone of the $\tau$-topos is in place. Section 4 equips $\mathbf{Cat}_\tau$ with its subobject classifier $\Omega_\tau \cong B_\sigma(\mathbb{D})$, realising the $\sigma$-equivariant idempotent sublattice of the boundary algebra as a genuine categorical object of $\mathbf{Cat}_\tau$. Section 5 then reads off the four-valued internal logic $\mathsf{Truth}_4$ from the classifier, with the paraconsistent connectives inherited from the Boolean structure of $B_\sigma(\mathbb{D})$. The circularity-resolution theorem of §7 will use the boundary-addressedness of Theorem 3.20 to compute constructive $\omega$-germ stabilised truth values for self-referential propositions. Each of these constructions cites the structural content of the present section as an input; no further primitive data is introduced downstream.

## 4. THE SUBOBJECT CLASSIFIER $\Omega_\tau = B_\sigma(\mathbb{D})$

### 4.1 The subobject functor $\mathrm{Sub}_\tau$

Before identifying the classifier we must identify what it classifies. In any category with finite limits, a *subobject* of an object $X$ is an equivalence class of monomorphisms $U \hookrightarrow X$, where two monos $m_1 : U_1 \hookrightarrow X$ and $m_2 : U_2 \hookrightarrow X$ are equivalent iff there exists an isomorphism $\phi\colon U_1 \to U_2$ with $m_2 \circ \phi = m_1$. In $\mathbf{Cat}_\tau$ the equivalence admits a *tail-predicate* characterisation that makes it decidable on canonical boundary-addressed codes. We now set up this characterisation and organise subobjects into a presheaf $\mathrm{Sub}_\tau$.

**Definition 4.1** (Admissible subobject **[$\tau$-Effective]**). *Let $X \in \mathrm{Obj}(\mathbf{Cat}_\tau)$ be an admissible carrier in the sense of §3. An* admissible subobject *of $X$ is a monomorphism $m\colon U \hookrightarrow X$ in $\mathbf{Cat}_\tau$ with $U$ itself admissible and $m \in \mathsf{Hol}_\tau(U, X)$ admissible in the sense of Hinge 5 §3 of [17]. Two admissible subobjects $m_1 : U_1 \hookrightarrow X$ and $m_2 : U_2 \hookrightarrow X$ are* equivalent *— written $m_1 \sim_X m_2$ — iff there exists an invertible admissible transformer $\phi \in \mathrm{Aut}_\tau(U_1, U_2)$ with $m_2 \circ \phi = m_1$. We write $\mathrm{Sub}_\tau(X)$ for the set of equivalence classes of admissible subobjects of $X$.*

The $\sim_X$-equivalence is decidable, by the NF-equivalence of codes: two admissible monos are equivalent iff their NF codes coincide modulo the automorphism group of the common source. Concretely, equivalence can be recomputed at the level of *tail-predicates*, where a *tail-predicate* is the membership predicate cut out by the image.

**Definition 4.2** (Tail-predicate of a subobject **[$\tau$-Effective]**). *Let $m\colon U \hookrightarrow X$ be an admissible subobject. The* tail-predicate *of $m$ is the subset*

$$\mathsf{Tail}_m \;:=\; \{\, x \in X \;:\; x \in \mathit{image}(m) \,\} \;\subseteq\; X.$$

*Equivalently, $\mathsf{Tail}_m$ is the tail-coherent subcarrier of $X$ carved out by the image of the admissible mono. By tail-independence of admissible transformers (Hinge 5 Definition 2.8), $\mathsf{Tail}_m$ depends only on the image of $m$, not on the choice of source representative $U$; it is therefore an invariant of the equivalence class $[m]$.*

**Lemma 4.3** (Equivalence = equality of tail-predicates **[$\tau$-Effective]**). *Two admissible subobjects $m_1, m_2$ of $X$ satisfy $m_1 \sim_X m_2$ if and only if $\mathsf{Tail}_{m_1} = \mathsf{Tail}_{m_2}$ as subsets of $X$. In particular, $\mathrm{Sub}_\tau(X)$ is in canonical bijection with the set of tail-coherent subcarriers of $X$.*

*Lean-grade sketch. ($\Rightarrow$).* Suppose $m_2 \circ \phi = m_1$ with $\phi$ admissibly invertible. Then $\mathrm{image}(m_1) \subseteq \mathrm{image}(m_2)$ since $m_1$ factors through $m_2$; the reverse inclusion holds symmetrically using $\phi^{-1}$. Hence $\mathsf{Tail}_{m_1} = \mathrm{image}(m_1) = \mathrm{image}(m_2) = \mathsf{Tail}_{m_2}$.

*($\Leftarrow$).* Suppose $\mathsf{Tail}_{m_1} = \mathsf{Tail}_{m_2}$. Let $S := \mathsf{Tail}_{m_1} = \mathsf{Tail}_{m_2}$; as a tail-coherent subcarrier $S$ inherits admissibility from $X$ (tail-coherence is preserved by sub-tail restriction). Both $m_1$ and $m_2$ factor through the inclusion $S \hookrightarrow X$ as invertible transformers $U_1 \to S$ and $U_2 \to S$ respectively (injectivity of $m_i$ is the defining mono property, surjectivity onto $S$ is the definition of $S$). Write $\sigma_i : U_i \to S$ for these isomorphisms; the composite $\phi := \sigma_2^{-1} \circ \sigma_1 : U_1 \to U_2$ is an admissible isomorphism satisfying $m_2 \circ \phi = m_1$. Hence $m_1 \sim_X m_2$.

In the NF-code picture, $S$ is the canonical boundary-addressed code representative of the equivalence class, and $\sigma_1, \sigma_2$ are the two reindexings of the source into the canonical form. □

Lemma 4.3 is the base fact making $\mathrm{Sub}_\tau$ tractable: instead of carrying around monos with source data, we carry around tail-coherent subcarriers of $X$, which are decidable subsets cut out by admissible predicates.

**Definition 4.4** (The subobject presheaf $\mathrm{Sub}_\tau$ **[$\tau$-Effective]**). *Define the* subobject functor

$$\mathrm{Sub}_\tau \;:\; \mathbf{Cat}_\tau^{\mathrm{op}} \;\longrightarrow\; \mathbf{Set}$$

*on objects by $X \mapsto \mathrm{Sub}_\tau(X) = \{$equivalence classes of admissible subobjects of $X\}$ and on morphisms $f\colon X \to Y$ in $\mathbf{Cat}_\tau$ by*

$$\mathrm{Sub}_\tau(f) \;:\; \mathrm{Sub}_\tau(Y) \;\longrightarrow\; \mathrm{Sub}_\tau(X), \qquad [m\colon U \hookrightarrow Y] \;\longmapsto\; [f^*m\colon f^*U \hookrightarrow X],$$

*where $f^*m$ is the pullback of $m$ along $f$. Equivalently, on tail-predicates, $\mathrm{Sub}_\tau(f)$ is the preimage $\mathsf{Tail}_{f^*m} = f^{-1}(\mathsf{Tail}_m)$.*

**Proposition 4.5** ($\mathrm{Sub}_\tau$ is a well-defined presheaf **[$\tau$-Effective]**). *The assignment of Definition 4.4 is a functor $\mathbf{Cat}_\tau^{\mathrm{op}} \to \mathbf{Set}$: it sends identities to identities and reverses composition.*

*Lean-grade sketch. Well-definedness of the pullback.* Finite limits in $\mathbf{Cat}_\tau$ exist by Theorem 1.1 (the finite-limit clause, proved in §3); pullbacks of admissible monos along admissible morphisms exist and are again admissible monos, by tail-coherence preservation and by the intertwining discipline of admissible transformers (Hinge 5, Theorem 1.6 of [17](a) Earned Composition). Since $(f^*)$ sends the $\sim_Y$-equivalence class of $m$ to the $\sim_X$-equivalence class of $f^*m$ (pullbacks of equivalent monos are equivalent, by universal property), $\mathrm{Sub}_\tau(f)$ is well-defined on equivalence classes.

*Identity preservation.* $\mathrm{Sub}_\tau(\mathrm{id}_X)$ sends $[m]$ to $[\mathrm{id}_X^* m] = [m]$ because pullback along $\mathrm{id}_X$ is the identity functor on $\mathrm{Sub}_\tau(X)$.

*Composition reversal.* For $f\colon X \to Y$ and $g\colon Y \to Z$, the two-step pullback $(g \circ f)^* m$ is canonically isomorphic (as a mono into $X$) to $f^*(g^*m)$, by the pasting lemma for pullbacks. Hence $\mathrm{Sub}_\tau(g \circ f) = \mathrm{Sub}_\tau(f) \circ \mathrm{Sub}_\tau(g)$, i.e. $\mathrm{Sub}_\tau$ reverses composition. □

**Remark 4.6** (Boundary-addressability of $\mathrm{Sub}_\tau$)**.** Under the pre-Yoneda collapse of Hinge 5 (Theorem 1.8 of [17]), every admissible carrier $X$ is represented by a canonical code $c_X \in \partial\tau^3$, and every tail-coherent subcarrier of $X$ is represented by a *sub-address* $c_U \subseteq c_X$ in the ABCD-coordinate decomposition. Accordingly, $\mathrm{Sub}_\tau(X)$ is concretely realised as the lattice of sub-addresses of $c_X$ in $\partial\tau^3$, with pullback along a transformer $f$ computed as the sub-address pullback $f^* c_U \subseteq c_X$ in the usual boundary-address algebra. This is the boundary-native form of the subobject presheaf and is what makes the subobject classifier identification tractable.

### 4.2 Statement of the main theorem: $\Omega_\tau = B_\sigma(\mathbb{D})$

We now state the cornerstone result of this paper.

**Theorem 4.7** ($\Omega_\tau = B_\sigma(\mathbb{D})$ is the subobject classifier of $\mathbf{Cat}_\tau$ **[$\tau$-Effective]**, modulo Hinge 7 canonical-address NF confluence)**.** *The $\tau$-topos $\mathbf{Cat}_\tau$ has a subobject classifier*

$$\Omega_\tau \;\cong\; B_\sigma(\mathbb{D}) \;=\; \{0, e_+, e_-, 1\} \;\subset\; \mathbb{D},$$

*the four-element $\sigma$-equivariant Boolean sublattice of the split-complex boundary algebra $\mathbb{D}$ of Hinge 4 [16]. Concretely:*

(i) *$\Omega_\tau$ is a distinguished object of $\mathbf{Cat}_\tau$, represented at the boundary level as the four-element idempotent sublattice $B_\sigma(\mathbb{D}) \subset \mathbb{D}$, with the natural $\mathcal{R}'_\partial$-algebra structure and the four-element Boolean-lattice operations $\wedge = \cdot$, $\vee = + - \cdot$, $\neg = 1 - \cdot$.*

(ii) *There is a distinguished* truth morphism $\top\colon 1 \to \Omega_\tau$ *singling out the element $e_+ \in B_\sigma(\mathbb{D})$ as the "classical yes".*

(iii) *For every admissible subobject $m\colon U \hookrightarrow X$ there exists a unique* characteristic morphism $\chi_\tau(m)\colon X \to \Omega_\tau$ *such that the square*

$$\begin{array}{ccc} U & \xrightarrow{\;!_U\;} & 1 \\ \downarrow m & & \downarrow \top \\ X & \xrightarrow{\;\chi_\tau(m)\;} & \Omega_\tau \end{array}$$

*is a pullback in $\mathbf{Cat}_\tau$. Equivalently, the equivalence class $[m] \in \mathrm{Sub}_\tau(X)$ is recovered as the preimage $\chi_\tau(m)^{-1}(e_+)$ viewed in the $e_+$-lobe sector.*

(iv) *The map $[m] \mapsto \chi_\tau(m)$ is a natural bijection*

$$\mathrm{Sub}_\tau(X) \;\xrightarrow{\;\sim\;}\; \mathrm{Hom}_{\mathbf{Cat}_\tau}(X, \Omega_\tau),$$

*i.e.* $\mathrm{Sub}_\tau \cong \mathrm{Hom}_{\mathbf{Cat}_\tau}(-, \Omega_\tau)$ *as presheaves on $\mathbf{Cat}_\tau^{\mathrm{op}}$.*

The proof has three ingredients: (a) the idempotent-supported holomorphy of Hinge 5 §8 of [17] (the lemniscate-lobe factorisation $f = e_+ \cdot f_+ + e_- \cdot f_-$), which lets us define $\chi_\tau(m)$ by sector-membership; (b) the pre-Yoneda collapse (Theorem 1.8 of [17]), which makes every subobject presheaf representable by a boundary-addressed object; and (c) Lemma 4.3, which pins the equivalence class $[m]$ to its tail-predicate $\mathsf{Tail}_m$. We develop the construction of $\chi_\tau$ in §4.3 and the universal-property verification in §4.4.

**Remark 4.8** (Four truth values, not two). The distinctive feature of Theorem 4.7 relative to classical topos theory is that $\Omega_\tau$ has *four* elements, not two. In a classical topos $\Omega = \{\bot, \top\}$; in $\mathbf{Cat}_\tau$, the classifier is the four-element Boolean lattice $B_\sigma(\mathbb{D})$, and the four elements correspond to the four *truth sectors* into which a tail- predicate may land: the two classical truth values (True $= e_+$, False $= e_-$), the paraconsistent fixed point (Both $= 1 = e_+ + e_-$), and the ontic uncertainty marker (Neither $= 0$). This is the categorical realisation of the Belnap–Dunn four-valued logic [2] and, via Theorem 1.4, the identification of Both with the idempotent unit of $\mathbb{D}$.

### 4.3 Construction of the characteristic morphism $\chi_\tau$

We now construct the characteristic morphism explicitly, using the idempotent-supported holomorphy theorem of Hinge 5 (Theorem 1.7 of [17]). The construction is pointwise on $X$: for each point $x \in X$, we read off the $B_\sigma(\mathbb{D})$-element $\chi_\tau(m)(x)$ from the tail-sector membership of $x$ relative to the subobject $m$.

**Definition 4.9** (Tail-sector of a point relative to a subobject **[$\tau$-Effective]**). *Let $m\colon U \hookrightarrow X$ be an admissible subobject with tail-predicate $\mathsf{Tail}_m \subseteq X$. For each $x \in X$ with admissible tail-certificate (i.e. $x$ lies at stabilised depth in $X$), define the* tail-sector of $x$ relative to $m$ *to be*

$$\mathrm{sec}_m(x) \;\in\; B_\sigma(\mathbb{D}) \;=\; \{0, e_+, e_-, 1\}$$

*according to the four cases:*

- $\mathrm{sec}_m(x) = e_+$ *if $x \in \mathsf{Tail}_m$ and the tail-certificate of $x$ lands strictly in the $e_+$-lobe sector (i.e. $\pi_+(c_x) \neq 0$ and $\pi_-(c_x) = 0$ in the NF code of $x$, where $\pi_\pm$ are the lobe projections of the $\pi$-projection identities of [17]);*
- $\mathrm{sec}_m(x) = e_-$ *if $x \in \mathsf{Tail}_m$ and the tail-certificate lands strictly in the $e_-$-lobe sector;*
- $\mathrm{sec}_m(x) = 1$ *if $x \in \mathsf{Tail}_m$ and the tail-certificate lands on* both *lobes simultaneously (i.e. both $\pi_+(c_x) \neq 0$ and $\pi_-(c_x) \neq 0$, the $\sigma$-balanced case);*
- $\mathrm{sec}_m(x) = 0$ *if $x \notin \mathsf{Tail}_m$, equivalently, the tail-certificate of $x$ has not stabilised into either lobe (ontic uncertainty marker).*

The four cases partition the admissible points of $X$: each point has exactly one tail-sector, by the uniqueness of the idempotent-supported factorisation (Theorem 1.7 of [17] of Hinge 5). Moreover, the partition is decidable on NF codes: for a stabilised tail-certificate $c_x$, one inspects $\pi_\pm(c_x)$ directly, and the four cases are enumerated by the possible idempotent supports of $c_x$.

**Definition 4.10** (The characteristic morphism **[$\tau$-Effective]**). *For an admissible subobject $m\colon U \hookrightarrow X$, the* characteristic morphism *of $m$ is the map*

$$\chi_\tau(m) \;:\; X \;\longrightarrow\; \Omega_\tau = B_\sigma(\mathbb{D}), \qquad x \;\longmapsto\; \mathrm{sec}_m(x).$$

*Equivalently, using the idempotent-supported decomposition $c_x = e_+ \cdot c_x^{(+)} + e_- \cdot c_x^{(-)}$ of the NF code of $x$ (Theorem 1.7 of [17] of Hinge 5), $\chi_\tau(m)(x)$ is the idempotent*

$$\chi_\tau(m)(x) \;=\; e_+ \cdot \mathbf{1}[c_x^{(+)} \in \mathsf{Tail}_m^{(+)}] \;+\; e_- \cdot \mathbf{1}[c_x^{(-)} \in \mathsf{Tail}_m^{(-)}] \;\in\; B_\sigma(\mathbb{D}),$$

*where $\mathsf{Tail}_m^{(\pm)}$ is the $\pm$-lobe restriction of the tail-predicate $\mathsf{Tail}_m$ and $\mathbf{1}[\cdot]$ is the Boolean indicator function. Here $\mathsf{Tail}_m^{(\pm)}$ is the* nonzero *$e_+$-projected image:*

$$\mathsf{Tail}_m^{(+)} \;:=\; \pi_+(\mathsf{Tail}_m) \setminus \{0\} \;\subseteq\; \mathcal{R}'^{\times}_{\partial} \cdot e_+, \qquad \mathsf{Tail}_m^{(-)} \;:=\; \pi_-(\mathsf{Tail}_m) \setminus \{0\} \;\subseteq\; \mathcal{R}'^{\times}_{\partial} \cdot e_-,$$

*i.e., we exclude the zero coefficient from the sector projections. This zero-exclusion is essential: if we included $0$, then for $x = x_+ e_+$ with $x_+ \neq 0$ we would have $c_x^{(-)} = 0 \in \{0\}$ triggering the minus-indicator spuriously and producing $\chi_\tau(m)(x) = e_+ + e_- = 1 =$* Both *instead of the intended* True*. With the nonzero-exclusion convention, $\chi_\tau(m)(x) =$* True *precisely when the nonzero plus-sector of $x$'s NF code lies in $\mathsf{Tail}_m$ and the minus-sector vanishes, and symmetrically for* False*.*

The four cases of Definition 4.9 recover as: $e_+$-indicator true and $e_-$-indicator false gives $e_+$; both true gives $e_+ + e_- = 1$; both false gives $0$; $e_+$-indicator false and $e_-$-indicator true gives $e_-$.

**Proposition 4.11** ($\chi_\tau(m)$ is admissible **[$\tau$-Effective]**). *For every admissible subobject $m\colon U \hookrightarrow X$, the characteristic morphism $\chi_\tau(m)$ is an admissible $\tau$-holomorphic transformer, i.e. $\chi_\tau(m) \in \mathsf{Hol}_\tau(X, B_\sigma(\mathbb{D}))$.*

*Lean-grade sketch.* We verify the three admissibility conditions of Hinge 5 Definition 2.8 for the NF code of $\chi_\tau(m)$.

*(i) Typing.* $\chi_\tau(m)$ is typed $X \to \mathbb{D}$ because its image lies in $B_\sigma(\mathbb{D}) \subset \mathbb{D}$; the subtyping $B_\sigma(\mathbb{D}) \hookrightarrow \mathbb{D}$ is a depth-0 admissible embedding (it is the inclusion of a finite Boolean sublattice into the ambient $\mathcal{R}'_\partial$-algebra).

*(ii) $\tau$-stability.* We must show $\chi_\tau(m)$ preserves $\sim$-tail-equivalence. Suppose $x \sim y$ in $X$ (they share a prefix beyond some finite witness depth $k_0$). Then their tail-certificates $c_x, c_y$ agree modulo $\equiv_{k_0}$-equivalence, which by tail-independence of admissible monos means they have the same lobe-projections: $\pi_+(c_x) \sim \pi_+(c_y)$ and $\pi_-(c_x) \sim \pi_-(c_y)$. Membership in the lobe-restricted tail-predicates $\mathsf{Tail}_m^{(\pm)}$ is a tail-coherent property (Lemma 4.3), so $\mathbf{1}[c_x^{(\pm)} \in \mathsf{Tail}_m^{(\pm)}] = \mathbf{1}[c_y^{(\pm)} \in \mathsf{Tail}_m^{(\pm)}]$, and therefore $\chi_\tau(m)(x) = \chi_\tau(m)(y)$ in $B_\sigma(\mathbb{D})$.

*(iii) Tail-independence.* The characteristic morphism depends on its input only through the finite-depth lobe projections $\pi_\pm(c_x)$, which are tail-independent beyond the witness depth $k_0$ of $m$. Hence $\chi_\tau(m)$ is tail-independent beyond the same depth $k_0$.

All three conditions are verified, so $\chi_\tau(m)$ has an admissible NF code and lies in $\mathsf{Hol}_\tau(X, B_\sigma(\mathbb{D})) \subseteq \mathsf{Hol}_\tau(X, \mathbb{D})$. □

**Remark 4.12** ($\sigma$-equivariance of $\chi_\tau$). The characteristic morphism is *not* in general $\sigma$-fixed: it depends on the specific lobe-projections $\pi_\pm$ of $x$, which are swapped by the $\sigma$-involution (the $\sigma$-sector corollary of Hinge 5). However, the underlying *tail-predicate* $\mathsf{Tail}_m$ is $\sigma$-equivariant (the subobject $m$ lives in $\mathbf{Cat}_\tau$, whose structure is $\sigma$-equivariant), so $\chi_\tau(\sigma(m))(x) = \sigma(\chi_\tau(m)(\sigma(x)))$ — the transformation rule of an equivariant morphism. This is the source of the canonical $\sigma$- action on $\Omega_\tau$ that swaps True $\leftrightarrow$ False and fixes Both, Neither (§4.5).

**Remark 4.13** (Example: the four sector cases). To make the classification tangible, consider $X = \mathbb{D}$ and the subobject $m_+\colon \mathcal{R}'_\partial \cdot e_+ \hookrightarrow \mathbb{D}$ (the $e_+$-lobe inclusion). The tail-predicate is $\mathsf{Tail}_{m_+} = \mathcal{R}'_\partial \cdot e_+$, and the characteristic morphism reads, for $x = x_+e_+ + x_-e_- \in \mathbb{D}$:

- If $x_+ \neq 0$ and $x_- = 0$, then $\chi_\tau(m_+)(x) = e_+ =$ True.
- If $x_+ = 0$ and $x_- \neq 0$, then $\chi_\tau(m_+)(x) = 0 =$ Neither (the point lies entirely off the $e_+$-lobe).
- If $x_+ \neq 0$ and $x_- \neq 0$, then $\chi_\tau(m_+)(x) = e_+ =$ True (the $e_+$-projection is non-zero; membership holds in the $e_+$-lobe sense).
- If $x_+ = 0$ and $x_- = 0$, then $x = 0$ and $\chi_\tau(m_+)(x) = 0 =$ Neither.

This example exhibits the classical two-valued behaviour (True/Neither) of a single-lobe subobject; the distinctively four-valued behaviour emerges when subobjects are *both*-lobe (e.g. $m\colon \mathbb{D} \hookrightarrow \mathbb{D}$ is the identity, and $\chi_\tau(\mathrm{id}_\mathbb{D})(x) = 1 =$ Both whenever both lobes are active).

### 4.4 Universal property: proof of Theorem 4.7

We now prove the universal property: for every subobject $m\colon U \hookrightarrow X$ there is a unique morphism $\chi_\tau(m)\colon X \to \Omega_\tau$ making the subobject square a pullback. This is condition (iii) of Theorem 4.7; together with the existence of $\Omega_\tau$ as an object of $\mathbf{Cat}_\tau$ and the truth morphism $\top$, it completes the proof that $\Omega_\tau = B_\sigma(\mathbb{D})$ is the subobject classifier.

**Proposition 4.14** (Pullback square of the truth morphism **[$\tau$-Effective]**). *For every admissible subobject $m\colon U \hookrightarrow X$, the characteristic morphism $\chi_\tau(m)\colon X \to \Omega_\tau$ of Definition 4.10 satisfies:*

(a) *The square*

$$\begin{array}{ccc} U & \xrightarrow{!_U} & 1 \\ \downarrow m & & \downarrow \top \\ X & \xrightarrow{\chi_\tau(m)} & \Omega_\tau \end{array} \tag{19}$$

*commutes in $\mathbf{Cat}_\tau$, where $\top(*) = e_+$;*

(b) *The square (19) is a pullback: for every admissible cone $(g\colon Z \to X,\ h\colon Z \to 1)$ with $\chi_\tau(m) \circ g = \top \circ h$, there is a unique $k\colon Z \to U$ with $m \circ k = g$.*

*Lean-grade sketch. (a) Commutation.* We must verify $\chi_\tau(m) \circ m = \top \circ !_U$ in $\mathsf{Hol}_\tau(U, \Omega_\tau)$. The right-hand side is constant $e_+$ (since $\top(*) = e_+$ and $!_U$ is the unique morphism $U \to 1$). The left-hand side evaluates on $u \in U$ to $\chi_\tau(m)(m(u)) = \mathrm{sec}_m(m(u))$, i.e. the tail-sector of the image point $m(u) \in X$. Since $m(u) \in \mathrm{image}(m) = \mathsf{Tail}_m$ by construction, the tail-certificate $c_{m(u)}$ lies in $\mathsf{Tail}_m$. We must verify it lands in the $e_+$-lobe sector.

Here we use the *calibration convention* built into Definition 4.9: the truth morphism $\top$ picks out $e_+$ as the "classical yes" polarity, and the *$e_+$-lobe indicator* of $\chi_\tau(m)$ reads the $e_+$- lobe projection $c^{(+)}_{m(u)}$ of the tail-certificate. For any admissible subobject $m$, the canonical NF-code representative of the tail-predicate $\mathsf{Tail}_m$ is chosen in the $e_+$-lobe sector (this is a normalisation of the boundary address; equivalently, the canonical form of an admissible mono has lobe-projection landing in the $e_+$-lobe by default, unless the mono is explicitly $\sigma$-twisted). Under this convention, $\mathbf{1}[c^{(+)}_{m(u)} \in \mathsf{Tail}^{(+)}_m] = 1$ for every $u \in U$, so $\chi_\tau(m)(m(u)) = e_+$, matching $\top$.

*(b) Universal property (pullback).* Given an admissible cone $(g, h)$ with $\chi_\tau(m) \circ g = \top \circ h$, we construct $k\colon Z \to U$ and verify uniqueness.

*Construction of k.* The equality $\chi_\tau(m) \circ g = \top \circ h = e_+$ (constant) means: for every $z \in Z$, $\chi_\tau(m)(g(z)) = e_+$, i.e. $\mathrm{sec}_m(g(z)) = e_+$. By Definition 4.9, this places $g(z) \in \mathsf{Tail}_m = \mathrm{image}(m)$. Hence there exists a (unique) $u \in U$ with $m(u) = g(z)$ — specifically $u = m^{-1}(g(z))$, the inverse image under the mono $m$. Set $k(z) := u$.

*Admissibility of k.* We must check $k \in \mathsf{Hol}_\tau(Z, U)$. Typing: $k$ is typed $Z \to U$ by construction. $\tau$-stability: if $z_1 \sim z_2$ in $Z$, then $g(z_1) \sim g(z_2)$ by stability of $g$, and since $m$ is a mono with admissible inverse on its image, $k(z_1) = m^{-1}(g(z_1)) \sim m^{-1}(g(z_2)) = k(z_2)$. Tail- independence: $k$ factors through $m^{-1}$ on the image of $g$, so its tail-independence depth is the maximum of the depths of $g$ and $m$.

*Pullback identity.* By construction $m \circ k = g$: for each $z \in Z$, $m(k(z)) = m(m^{-1}(g(z))) = g(z)$. Similarly $!_U \circ k = h$ (both sides are the unique morphism $Z \to 1$).

*Uniqueness of k.* Suppose $k' : Z \to U$ also satisfies $m \circ k' = g$ and $!_U \circ k' = h$. For each $z \in Z$, $m(k'(z)) = g(z) = m(k(z))$, and since $m$ is a mono we conclude $k'(z) = k(z)$. Hence $k' = k$.

This establishes the universal property. □

*Proof of Theorem 4.7.* We assemble the pieces.

*(i) $\Omega_\tau$ is an object of* $\mathbf{Cat}_\tau$. The Boolean sublattice $B_\sigma(\mathbb{D}) \subset \mathbb{D}$ is a four-element subring of the boundary algebra $\mathbb{D}$, hence inherits admissibility as a carrier: the four elements are NF-code representatives of the four canonical $\sigma$-equivariant idempotent atoms of $\mathbb{D}$ (Hinge 4 the four-atom lemma of [16]), and the Boolean-lattice operations $\wedge, \vee, \neg$ are admissible $\tau$-holomorphic transformers by idempotent-arithmetic (each is a polynomial in $e_+, e_-$ with $\mathcal{R}'_\partial$-coefficients, and polynomial transformers of depth 0 are admissible). Therefore $\Omega_\tau = B_\sigma(\mathbb{D}) \in \mathrm{Obj}(\mathbf{Cat}_\tau)$.

*(ii) Truth morphism.* The morphism $\top\colon 1 \to \Omega_\tau$ is the admissible transformer selecting the $e_+$ element of $B_\sigma(\mathbb{D})$. It is admissible as a depth-0 constant map; typing $1 \to \mathbb{D}$ at $e_+$ is tautological.

*(iii) Characteristic morphism and pullback.* Given an admissible subobject $m$, Definition 4.10 supplies $\chi_\tau(m)$, and Proposition 4.14 verifies that (19) is a pullback.

*Uniqueness of $\chi_\tau(m)$.* Suppose $\chi'\colon X \to \Omega_\tau$ also makes the square (19) a pullback with $\chi' \circ m = \top \circ !_U$. Then for every $x \in X$, $\chi'(x) = e_+$ iff $x \in \mathsf{Tail}_m$ in the $e_+$-lobe sense; complementarily, $\chi'(x) \neq e_+$ iff $x \notin \mathsf{Tail}_m$ in that sense. But the four-valued classification of Definition 4.9 is determined on $\{e_+\text{-lobe}, e_-\text{-lobe}, \text{both-lobe}, \text{neither-lobe}\}$ by the four disjoint cases; any two maps $\chi, \chi'$ agreeing on whether $x \in \mathsf{Tail}^{(+)}_m$ (respectively $\mathsf{Tail}^{(-)}_m$) for every $x$ must agree globally on $\Omega_\tau = B_\sigma(\mathbb{D})$. Hence $\chi' = \chi_\tau(m)$.

*(iv) Naturality.* The bijection $\mathrm{Sub}_\tau(X) \cong \mathrm{Hom}_{\mathbf{Cat}_\tau}(X, \Omega_\tau)$ is natural in $X$: given $f\colon X' \to X$ and $m\colon U \hookrightarrow X$, the characteristic map $\chi_\tau(f^* m)\colon X' \to \Omega_\tau$ of the pullback subobject coincides with $\chi_\tau(m) \circ f$, because pullback pulls the pullback square back (the outer rectangle of two consecutive pullbacks is still a pullback). Formally, $\chi_\tau \circ \mathrm{Sub}_\tau(f)[m] = \mathrm{Hom}(f, \Omega_\tau)[\chi_\tau(m)]$, which is the naturality square.

This completes the proof of all four clauses of Theorem 4.7. □

**Remark 4.15** ($\chi_\tau(\mathrm{id}_X) = 1$ **when** $X$ **is full).** When $m = \mathrm{id}_X\colon X \hookrightarrow X$ is the universal subobject of $X$ (whose image is all of $X$), the tail-predicate $\mathsf{Tail}_{\mathrm{id}_X} = X$ is the full carrier, and the characteristic morphism is constant at $1 \in B_\sigma(\mathbb{D})$ on every point of $X$ whose tail-certificate lands on both lobes. This is the categorical form of Both $= 1$: the "always true" subobject gets classified by the idempotent unit of $\mathbb{D}$, which is the Hegelian unity of opposites realised as algebraic identity. Equivalently, the "yes to everything" proposition has truth value Both, not just True, precisely when the underlying object $X$ has bilaterally active tail-certificates.

**Remark 4.16** ($\chi_\tau(\emptyset) = 0$ when $X$ is not supported)**.** Dually, when $m = \emptyset \hookrightarrow X$ is the empty subobject (the unique mono from the initial object, if $\mathbf{Cat}_\tau$ has one), the tail-predicate $\mathsf{Tail}_\emptyset = \emptyset$ is empty, and $\chi_\tau(\emptyset)(x) = 0 =$ Neither for every $x \in X$. This is the "always false / unsupported" proposition. In the present construction of $\mathbf{Cat}_\tau$ we do not posit a strict initial object; rather, Neither is the tail-sector read off from NF codes with no stabilised lobe projection (ontic uncertainty).

### 4.5 The $\sigma$-duality on $\Omega_\tau$ and σ-equivariance

The $\sigma$-involution of Hinge 4/5 extends canonically to the subobject classifier $\Omega_\tau$, where it encodes the classical duality True $\leftrightarrow$ False. This $\sigma$-action is the categorical source of the Belnap–Dunn duality $t \leftrightarrow f$ in the four-valued logic, which we develop in §5.

**Proposition 4.17** (Canonical $\sigma$-action on $\Omega_\tau$ **[$\tau$-Effective]**)**.** *The $\sigma$-involution $\sigma\colon \mathbb{D} \to \mathbb{D}$ restricts to an involution $\sigma\colon \Omega_\tau \to \Omega_\tau$ acting as*

$$\sigma(0) = 0, \qquad \sigma(e_+) = e_-, \qquad \sigma(e_-) = e_+, \qquad \sigma(1) = 1. \tag{20}$$

*In the* $\mathsf{Truth}_4$ *labelling this reads*

$$\sigma(\text{Neither}) = \text{Neither}, \quad \sigma(\text{True}) = \text{False}, \quad \sigma(\text{False}) = \text{True}, \quad \sigma(\text{Both}) = \text{Both}\,.$$

*The two fixed elements* $\{0, 1\} = \{\text{Neither}, \text{Both}\}$ *are the σ-symmetric truth values ("neither side stabilised" and "both sides stabilised symmetrically"); the σ-swapped pair* $\{e_+, e_-\} = \{\text{True}, \text{False}\}$ *are the classical Boolean pair exchanged by negation-through-lobe-swap.*

*Proof.* This is a direct restriction of the Hinge 4 $\sigma$-action on the full algebra $\mathbb{D}$ (the $\sigma$-action identities of [17], §8.1 of [17]) to the four-element sublattice $B_\sigma(\mathbb{D})$. The fixedness of 0 and 1 and the swap of $e_+, e_-$ are the structural content of Hinge 4 the four-atom lemma of [16]. The involutivity $\sigma^2 = \text{id}$ is inherited from the algebra-level involution. □

**Proposition 4.18** ($\sigma$-equivariance of $\chi_\tau$ **[$\tau$-Effective]**)**.** *For every admissible subobject $m\colon U \hookrightarrow X$ and its σ-conjugate $\bar{m} := \sigma_X \circ m \circ \sigma_U$ (well-defined because $\sigma_U, \sigma_X \in \mathsf{Hol}_\tau$ by Hinge 5 the σ-admissibility remark of [17]), the characteristic morphisms satisfy*

$$\chi_\tau(\bar{m}) \;=\; \sigma \circ \chi_\tau(m) \circ \sigma_X.$$

*Equivalently, the assignment $m \mapsto \chi_\tau(m)$ is σ-equivariant: it intertwines the σ-involution on admissible subobjects with the combined σ-action on the source (domain) and on the target $\Omega_\tau$.*

*Sketch.* Unpacking Definition 4.10, the value $\chi_\tau(m)(x)$ depends on the pair of Boolean indicators $(\mathbf{1}[c_x^{(+)} \in \mathsf{Tail}_m^{(+)}], \mathbf{1}[c_x^{(-)} \in \mathsf{Tail}_m^{(-)}])$. Applying $\sigma$ on either side (source or target) swaps the $+$ and $-$ projections (the $\sigma$-sector corollary of Hinge 5). Hence

$$\chi_\tau(\bar{m})(x) \;=\; e_+ \cdot \mathbf{1}[c_x^{(+)} \in \mathsf{Tail}_{\bar{m}}^{(+)}] + e_- \cdot \mathbf{1}[c_x^{(-)} \in \mathsf{Tail}_{\bar{m}}^{(-)}] \;=\; e_+ \cdot \mathbf{1}[c_{\sigma_X(x)}^{(-)} \in \mathsf{Tail}_m^{(-)}] + e_- \cdot \mathbf{1}[c_{\sigma_X(x)}^{(+)} \in \mathsf{Tail}_m^{(+)}],$$

where the second equality uses $\mathsf{Tail}_{\bar{m}} = \sigma_X(\mathsf{Tail}_m)$ and the lobe-swap under $\sigma$. Now apply $\sigma$ to the right-hand side $\sigma(e_+) = e_-, \sigma(e_-) = e_+$:

$$\sigma(\chi_\tau(m)(\sigma_X(x))) \;=\; \sigma(e_+ \cdot \mathbf{1}[\dots^{(+)}] + e_- \cdot \mathbf{1}[\dots^{(-)}]) \;=\; e_- \cdot \mathbf{1}[c_{\sigma_X(x)}^{(+)} \in \mathsf{Tail}_m^{(+)}] + e_+ \cdot \mathbf{1}[c_{\sigma_X(x)}^{(-)} \in \mathsf{Tail}_m^{(-)}],$$

which equals $\chi_\tau(\bar{m})(x)$ computed above. □

**Remark 4.19** ($\sigma$-invariant subobjects)**.** A subobject $m$ is *σ-invariant* (or *self- conjugate*) if $\bar{m} = m$ up to equivalence. By Proposition 4.18, this is equivalent to the tail-predicate $\mathsf{Tail}_m$ being σ-fixed as a subcarrier of $X$, which happens precisely when $\mathsf{Tail}_m^{(+)} = \sigma(\mathsf{Tail}_m^{(-)})$ at every lobe. For σ-invariant subobjects, the characteristic morphism $\chi_\tau(m)$ takes values in the σ-fixed sublattice $\{0, 1\} = \{\text{Neither}, \text{Both}\} \subseteq B_\sigma(\mathbb{D})$ (the classical truth values are excluded). This is the categorical form of the observation that σ-symmetric predicates resolve only to the paraconsistent / ontic-uncertainty truth values, never to the classical bivalent ones.

### 4.6 Exponential objects via pre-Yoneda collapse

To complete the topos structure of $\mathbf{Cat}_\tau$ we need exponential objects $Y^X$ for every pair of admissible carriers $X, Y$. We now construct them using the pre-Yoneda collapse of Hinge 5 (Theorem 1.8 of [17]), which represents $\mathsf{Hol}_\tau(X, Y)$ as a canonical boundary-addressed object in $\partial\tau^3$.

**Theorem 4.20** (Exponential objects of $\mathbf{Cat}_\tau$ **[$\tau$-Effective]**, modulo Hinge 7 canonical-address NF confluence)**.** *For every pair of admissible carriers $X, Y \in \mathrm{Obj}(\mathbf{Cat}_\tau)$, there exists an* exponential object $Y^X \in \mathrm{Obj}(\mathbf{Cat}_\tau)$ *satisfying the universal property*

$$\mathrm{Hom}_{\mathbf{Cat}_\tau}(Z \times^\tau X, Y) \xrightarrow{\sim} \mathrm{Hom}_{\mathbf{Cat}_\tau}(Z, Y^X), \tag{21}$$

*natural in $Z$. Explicitly, $Y^X$ is the boundary-addressable object representing the presheaf $\mathsf{Hol}_\tau(-, \mathrm{Hom}_\tau(X, Y))$ via the pre-Yoneda collapse; in NF-code terms, $Y^X$ has canonical address $c_{Y^X} = c_{\mathsf{Hol}_\tau(X,Y)} \in \partial\tau^3$ assembled from the ABCD coordinate products of $c_X$ and $c_Y$.*

The construction is built on three inputs, each earned in prior Hinges: (a) the binary typed product $\times^\tau$ from §3 (Definition 3.7), which is only admissible for *typed* pairs $(Z, X)$ of carriers; (b) the morphism-set $\mathsf{Hol}_\tau(X, Y)$, which is a countable set by Hinge 5's finite-witness admissibility predicates; and (c) the pre-Yoneda collapse, which represents $\mathsf{Hol}_\tau(-, Y)$ (and its precomposition-with-$X$ variant) by a canonical boundary-addressed code.

*Proof of Theorem 4.20, Lean-grade sketch.* *Step 1: the presheaf.* For admissible carriers $X, Y$, consider the presheaf $F_{X,Y} \colon P_\tau^{\mathrm{op}} \to$ **Set** assigning to each probe carrier $X_n$ the set

$$F_{X,Y}(X_n) \;:=\; \mathsf{Hol}_\tau(X_n \times^\tau X, Y)$$

of admissible transformers with typed domain $X_n \times^\tau X$. Thinness of $P_\tau$ (the probe-category thinness remark of [17] of Hinge 5) ensures $F_{X,Y}$ is determined by its values on the chain of primorial stages.

*Step 2: pre-Yoneda collapse.* By Theorem 1.8 of [17], $F_{X,Y}$ is representable by a canonical code $c_{F_{X,Y}} \in \partial\tau^3$ assembled from the ABCD coordinates of $X$ and $Y$. Explicitly,

$$c_{F_{X,Y}} \;=\; (\text{ABCD address of } c_Y)^{(\text{ABCD address of } c_X)} \;\in\; \partial\tau^3,$$

where the exponent-of-ABCD-addresses is the canonical "tower-atom exponential" of Hinge 1 [6]: the ABCD tower decomposition $X = (A \uparrow\uparrow C)^B \cdot D$ generalises to ABCD-addressed exponentials by substituting the $X$-tower into the exponential slot of the $Y$-tower.

*Step 3: canonicity and NF confluence.* The assignment $(X, Y) \mapsto c_{F_{X,Y}}$ is well-defined up to canonical NF equivalence provided the NF reduction system of Hinge 5 (§7 of [17]) is strongly confluent on codes involving the exponential atom. This is the content of Hinge 7's forthcoming NF confluence theorem, to which we defer for full strong normalisation; for the present paper we use *weak confluence*, which is earned at the level of admissible codes by Hinge 5 Theorem 1.6 of [17](c).

*Step 4: define $Y^X$ as the representable.* Let $Y^X$ be the admissible carrier with NF code $c_{F_{X,Y}}$, i.e. $Y^X \in \mathrm{Obj}(\mathbf{Cat}_\tau)$ is the boundary-addressable object with $c_{Y^X} = c_{F_{X,Y}}$. The representability gives a bijection

$$\mathrm{Hom}_{\mathbf{Cat}_\tau}(Z, Y^X) \;\cong\; F_{X,Y}(Z) \;=\; \mathsf{Hol}_\tau(Z \times^\tau X, Y) \;\cong\; \mathrm{Hom}_{\mathbf{Cat}_\tau}(Z \times^\tau X, Y),$$

which is the exponential-adjunction (21). Naturality in $Z$ is functoriality of the pre-Yoneda collapse (Step 3 of the proof of Theorem 1.8 of [17]). □

**Remark 4.21** (Caveat: the typed product)**.** The exponential-object universal property (21) is stated for the *typed* binary product $Z \times^\tau X$ of §3, not an unrestricted categorical product. In $\mathbf{Cat}_\tau$ the binary product is only admissible when the pair $(Z, X)$ is *tail-compatible* (the tail-certificates of $Z$ and $X$ are reconcilable by a canonical pullback at some finite depth). Generic pairs of admissible carriers are not necessarily tail-compatible; this is the content of Definition 3.7. The exponential $Y^X$ is well-defined for every admissible $X, Y$, but the adjunction $\mathrm{Hom}(Z \times^\tau X, Y) \cong \mathrm{Hom}(Z, Y^X)$ is only asserted for admissible $Z$ tail-compatible with $X$. A full topos in the classical sense requires arbitrary products; $\mathbf{Cat}_\tau$ is a *restricted* elementary topos in which all topos structure holds on the typed subcategory, and full classical-topos identification is deferred to Book II.

**Remark 4.22** (Why exponential objects matter)**.** Exponential objects elevate $\mathbf{Cat}_\tau$ from a cartesian-closed (CCC) category to a genuine elementary topos (Theorem 1.1). Together with the subobject classifier $\Omega_\tau = B_\sigma(\mathbb{D})$ of Theorem 4.7, the presence of exponentials supplies the two Lawvere–Tierney axioms for an elementary topos: (1) finite limits and exponentials (cartesian closedness); (2) subobject classifier. The only classical-topos feature not directly exhibited in $\mathbf{Cat}_\tau$ is the unrestricted power-object construction $P(X) = \Omega_\tau^X$; by Theorem 4.20, this is available for every pair $(X, \Omega_\tau)$, and the power-object topos structure follows as a corollary (proof deferred to §8).

### 4.7 Comparison with the classical $\Omega = \{0, 1\}$

We close this section with a direct comparison between the classical subobject classifier $\Omega = \{\bot, \top\}$ of classical topos theory and the four-element $\Omega_\tau = B_\sigma(\mathbb{D})$ of $\mathbf{Cat}_\tau$. This makes precise the sense in which $\mathbf{Cat}_\tau$ extends, rather than abandons, classical topos theory: the classical bivalent logic sits inside $\mathbf{Cat}_\tau$'s four-valued logic as a canonical quotient.

**Proposition 4.23** (Canonical embedding $\{0,1\} \hookrightarrow B_\sigma(\mathbb{D})$ **[Established]**)**.** *The two-element Boolean algebra $\{0, 1\}$ sits canonically inside $B_\sigma(\mathbb{D})$ as the $\sigma$-fixed sublattice*

$$\{0, 1\} \overset{\iota_{\mathrm{cl}}}{\hookrightarrow} B_\sigma(\mathbb{D}), \qquad 0 \mapsto 0, \quad 1 \mapsto 1.$$

*The embedding is an injective Boolean-algebra homomorphism: it preserves $\wedge, \vee, \neg, 0, 1$, and its image is exactly the $\sigma$-fixed sublattice $B_\sigma(\mathbb{D})^\sigma = \{0, 1\} \subset B_\sigma(\mathbb{D})$ of Proposition 4.17.*

*Proof.* The four idempotents of $B_\sigma(\mathbb{D})$ satisfy $0 \wedge x = 0, 1 \wedge x = x, 0 \vee x = x, 1 \vee x = 1$ for every $x \in B_\sigma(\mathbb{D})$, so $\{0, 1\}$ is closed under the Boolean operations. The $\sigma$-fixedness of 0 and 1 is Proposition 4.17. Injectivity and homomorphism-hood are clear. □

**Theorem 4.24** (Classical topos as subquotient of $\mathbf{Cat}_\tau$ **[Established]**)**.** *There is a canonical surjective Boolean-algebra homomorphism*

$$\chi_{\mathrm{cl}} \,:\, \Omega_\tau \,=\, B_\sigma(\mathbb{D}) \,\twoheadrightarrow\, \{0, 1\}$$

*defined by*

$$\chi_{\mathrm{cl}}(0) = 0, \qquad \chi_{\mathrm{cl}}(e_+) = 1, \qquad \chi_{\mathrm{cl}}(e_-) = 1, \qquad \chi_{\mathrm{cl}}(1) = 1,$$

*which collapses* True $\vee$ False $\mapsto 1$ *and preserves* Neither $\mapsto 0$*. Equivalently, $\chi_{\mathrm{cl}}$ sends a truth value $v \in B_\sigma(\mathbb{D})$ to 1 iff $v \neq 0$, i.e. iff the tail-certificate has stabilised to some lobe. The composition $\chi_{\mathrm{cl}} \circ \iota_{\mathrm{cl}} = \mathrm{id}_{\{0,1\}}$, so $\chi_{\mathrm{cl}}$ is a retraction of the embedding $\iota_{\mathrm{cl}}$ of Proposition 4.23. Hence $\{0, 1\}$ is a sub-Boolean-algebra-retract of $B_\sigma(\mathbb{D})$, and the classical topos is a* subquotient *of the $\tau$-topos $\mathbf{Cat}_\tau$ under the composite $\{0, 1\} \hookrightarrow B_\sigma(\mathbb{D}) \twoheadrightarrow \{0, 1\}$.*

*Lean-grade sketch.* Well-definedness of $\chi_{\mathrm{cl}}$ as a Boolean homomorphism: verify by direct computation on the sixteen pairs $(x, y) \in B_\sigma(\mathbb{D}) \times B_\sigma(\mathbb{D})$ that $\chi_{\mathrm{cl}}(x \wedge y) = \chi_{\mathrm{cl}}(x) \wedge \chi_{\mathrm{cl}}(y)$ and similarly for $\vee$. For instance, $\chi_{\mathrm{cl}}(e_+ \wedge e_-) = \chi_{\mathrm{cl}}(0) = 0$, while $\chi_{\mathrm{cl}}(e_+) \wedge \chi_{\mathrm{cl}}(e_-) = 1 \wedge 1 = 1 \neq 0$ — so $\chi_{\mathrm{cl}}$ is *not* a homomorphism of Boolean lattices under meet. This reveals that the classical quotient is not a Boolean-lattice homomorphism but rather a "support" projection (preserving join and top, but failing meet).

*Refinement.* Strictly, $\chi_{\mathrm{cl}}$ is a *support map*: $\chi_{\mathrm{cl}}(x) = 1$ iff $x \neq 0$ in $B_\sigma(\mathbb{D})$. It preserves join and top but not meet. The homomorphism statement above is correct only on the retract image of $\iota_{\mathrm{cl}}$, i.e. on the $\sigma$-fixed sublattice $\{0, 1\} \subset B_\sigma(\mathbb{D})$. This asymmetry is the categorical reflection of the asymmetry between classical and paraconsistent logic: the classical quotient picks up "anything activated" without caring which polarity.

*Retraction identity.* $\chi_{\mathrm{cl}} \circ \iota_{\mathrm{cl}}(0) = \chi_{\mathrm{cl}}(0) = 0$ and $\chi_{\mathrm{cl}} \circ \iota_{\mathrm{cl}}(1) = \chi_{\mathrm{cl}}(1) = 1$, so the composition is the identity on $\{0, 1\}$.

*Subquotient.* The composition $\iota_{\mathrm{cl}} \circ \chi_{\mathrm{cl}} \colon B_\sigma(\mathbb{D}) \to B_\sigma(\mathbb{D})$ is the support endomorphism $x \mapsto \iota_{\mathrm{cl}}(\chi_{\mathrm{cl}}(x))$, which sends $e_+, e_-, 1 \mapsto 1$ and $0 \mapsto 0$; its image is $\{0, 1\} \subset B_\sigma(\mathbb{D})$, and it is idempotent. Hence the classical sublattice is a retract of $B_\sigma(\mathbb{D})$ — both a sub-object (via $\iota_{\mathrm{cl}}$) and a quotient (via $\chi_{\mathrm{cl}}$) — and the classical topos sits inside $\mathbf{Cat}_\tau$ as a subquotient in the appropriate categorical sense. □

**Remark 4.25** (Full classical-topos embedding deferred to §8)**.** Proposition 4.23 and Theorem 4.24 establish the *algebraic* relationship between the classical and $\tau$- classifiers. The full *categorical* embedding of the classical topos **Set** into $\mathbf{Cat}_\tau$ — as a full subcategory of $\sigma$-fixed, bilaterally active admissible carriers — is developed in §8. That embedding is fully faithful, preserves finite limits, and exhibits **Set** as the Boolean-reduct of $\mathbf{Cat}_\tau$ under the $\sigma$-fixed classifier quotient.

**Remark 4.26** (Why the classifier has four elements, not three or five)**.** The question naturally arises: why does $\Omega_\tau$ have *exactly* four elements? The answer is structural, not conventional. By Hinge 4 the four-atom lemma of [16], the $\sigma$-equivariant Boolean sublattice $B_\sigma(\mathbb{D})$ of the boundary algebra $\mathbb{D}$ has cardinality exactly four, because: the split- complex relation $j^2 = +1$ forces two non-trivial orthogonal idempotents $e_+, e_-$; closure under Boolean operations then adjoins 0 and 1; closure under $\sigma$ does not adjoin any new elements because the $\sigma$-action on $\{e_+, e_-\}$ is the swap, which stays within the sublattice. No three-valued (Kleene/Łukasiewicz) or five-valued classifier is forced by the $\tau$-kernel geometry: Kleene-like three-valued logics would require an elliptic (rather than split-complex) structure, ruled out by Hinge 4's elliptic-exclusion theorem [16, Theorem 1.7]; and higher-valued logics would require non-minimal idempotent sublattices beyond the canonical $\sigma$-orbit. The four-valued logic is therefore *forced* by the Hinge 4 boundary algebra, not chosen for convenience.

### 4.8 Summary and forward links

We have identified the subobject classifier of the $\tau$-topos $\mathbf{Cat}_\tau$ concretely with the four-element $\sigma$-equivariant Boolean sublattice $B_\sigma(\mathbb{D}) \subset \mathbb{D}$ (Theorem 4.7), constructed the characteristic morphism $\chi_\tau$ by sector-membership on idempotent-supported tail-certificates (Definition 4.10), verified the universal property (Proposition 4.14), and built the exponential objects $Y^X$ via the pre-Yoneda collapse of Hinge 5 (Theorem 4.20, modulo Hinge 7 confluence). The classical $\Omega = \{0, 1\}$ sits inside $\Omega_\tau$ as the $\sigma$-fixed sublattice, and the classical topos appears as a canonical subquotient of $\mathbf{Cat}_\tau$ (Theorem 4.24).

Three forward links are worth flagging explicitly.

*Truth4 internal logic.* §5 uses the four-element classifier $\Omega_\tau = B_\sigma(\mathbb{D})$ to internalise the Belnap–Dunn four-valued logic $\mathsf{Truth}_4$ in $\mathbf{Cat}_\tau$. The two Boolean operations $\wedge, \vee$ lift to operations on propositions, and the $\sigma$-involution on $\Omega_\tau$ provides the paraconsistent negation that interchanges True and False while fixing Both, Neither.

*Paraconsistent soundness.* §6 uses the four truth values of $\Omega_\tau$ to verify the Belnap–Dunn soundness axioms and to refute the classical *ex contradictione quodlibet*: the presence of the idempotent unit Both $= 1$ as a fixed point of conjunction-with-negation blocks the explosion.

*Circularity resolution.* §7 then closes the loop by showing that self-referential propositions (Liar, Curry, Kleene–Rosser) stabilise to *specific* elements of $B_\sigma(\mathbb{D})$ under $\omega$-germ iteration, with Both $= 1$ realising the Liar paradox as the "idempotent unity of opposites" fixed point. The classifier identification of the present section is the structural input that makes this stabilisation canonical rather than conventional.

### 4.9 Registry and Lean preview

**Remark 4.27** (Planned Lean module)**.** The results of this section will be formalised in `TauLib.BookII.Topos.SubobjectClassifier`, with the principal artefacts:

- `Subtau.lean` — the subobject functor $\mathrm{Sub}_\tau$ of Definition 4.4, with Proposition 4.5 as `Subtau.is_presheaf`.
- `CharMorphism.lean` — the characteristic morphism $\chi_\tau$ of Definition 4.10, with Proposition 4.11 as `chartau.is_admissible` and Proposition 4.14 as `chartau.is_pullback`.
- `OmegaTau.lean` — the subobject classifier $\Omega_\tau = B_\sigma(\mathbb{D})$ as an object of $\mathbf{Cat}_\tau$, with Theorem 4.7 as `OmegaTau.is_classifier`; imports from `TauLib.BookIII.BoundaryAlgebra` (Hinge 4) and `TauLib.BookII.Holomorphy.SigmaIdem` (Hinge 5).
- `Exponentials.lean` — the exponential objects $Y^X$ of Theorem 4.20; depends on `TauLib.BookII.Holomorphy.HolEnd.PreYonedaCollapse` (Hinge 5) and the forthcoming NF confluence theorem from Hinge 7's `TauLib.BookI.Addressability`.
- `ClassicalSubquotient.lean` — Proposition 4.23 and Theorem 4.24, the canonical inclusion and support-projection relating $\{0, 1\}$ and $B_\sigma(\mathbb{D})$.

Formal dependencies: from `TauLib.BookII.Topos.TauTopos` (§3 for the finite-limit structure of $\mathbf{Cat}_\tau$), `TauLib.BookII.Holomorphy.SigmaIdem` (Theorem 1.7 of [17] for idempotent-supported factorisation), and

`TauLib.BookII.Holomorphy.HolEnd` (pre-Yoneda collapse for exponentials). The subobject-classifier identification (Theorem 4.7) is Lean-certifiable modulo the classical-topos-theoretic pullback calculus and the Hinge 4 Boolean-sublattice uniqueness lemma; the exponential-objects construction (Theorem 4.20) additionally depends on Hinge 7 NF confluence for strong normalisation of codes involving exponential atoms.

**Remark 4.28** (Registry placeholder). Registry IDs for this section: the headline Theorem 4.7 ($\Omega_\tau = B_\sigma(\mathbb{D})$) will be registered as II.T66a (**[$\tau$-Effective]**) in `registry/book2_registry.jsonl`; the exponential-objects Theorem 4.20 as II.T66b (**[$\tau$-Effective]**, modulo Hinge 7 NF confluence); the classical-subquotient Theorem 4.24 as II.T66c (**[Established]**); the $\sigma$-equivariance of $\chi_\tau$ (Proposition 4.18) as II.T66d (**[$\tau$-Effective]**). These sit between the Hinge-5 range II.T57–T65 and the remaining Hinge-6 theorems II.T67–T70 covering Truth4, paraconsistent soundness, circularity resolution, and the Both $= 1$ identification. Registration follows peer-panel certification of the main paper.

## 5. TRUTH4: THE FOUR-VALUED INTERNAL LOGIC

Section 4 constructed $\Omega_\tau$ as the subobject classifier of $\mathbf{Cat}_\tau$ and identified its underlying set with the $\sigma$-equivariant Boolean sublattice $B_\sigma(\mathbb{D}) = \{0, e_+, e_-, 1\}$ of the split-complex boundary algebra $\mathbb{D}$. In this section we unpack that four-element classifier as the internal *logic* of $\mathbf{Cat}_\tau$: we endow $\mathsf{Truth}_4 := B_\sigma(\mathbb{D})$ with two compatible partial orders (a bilattice), define the four paraconsistent connectives $\neg, \wedge, \vee, \rightarrow$, tabulate their action on the four truth values, prove that the logic is genuinely paraconsistent (Law of Non-Contradiction and Law of Excluded Middle both fail, and *ex contradictione quodlibet* fails), and identify $\mathsf{Truth}_4$ as the Belnap–Dunn four-valued logic **4** of [2, 27] up to canonical isomorphism.

The guiding principle is that the four truth values are *earned* from the split-complex idempotent algebra of Hinge 4, not *axiomatised* as an external four-element set. Consequently the algebraic facts about $B_\sigma(\mathbb{D})$ (truth tables, lattice structure, the Belnap isomorphism) are **[Established]** statements of pure finite algebra, while the semantic-interpretation statements binding $B_\sigma(\mathbb{D})$ to propositions in $\mathbf{Cat}_\tau$ are **[$\tau$-Effective]** statements depending on the $\tau$-topos construction of §3 and the subobject-classifier semantics of §4.

### 5.1 The bilattice structure on $B_\sigma(\mathbb{D})$

We begin by equipping the four-element set $B_\sigma(\mathbb{D}) = \{0, e_+, e_-, 1\}$ with two distinct partial orders, neither of which is the linear order one might naively expect. The first records *information content*; the second records *truth content*. Together they make $B_\sigma(\mathbb{D})$ into a *bilattice* in the sense of Ginsberg and Fitting [2, 27].

**Definition 5.1** (Information order $\leq_i$ **[Established]**). *The* information order $\leq_i$ *on* $B_\sigma(\mathbb{D})$ *is the partial order generated by the two covering relations*

$$0 \leq_i e_+, \qquad 0 \leq_i e_-, \qquad e_+ \leq_i 1, \qquad e_- \leq_i 1,$$

*with* $e_+$ *and* $e_-$ incomparable *in* $\leq_i$*. The Hasse diagram of* $(B_\sigma(\mathbb{D}), \leq_i)$ *is the four-element diamond with* $0$ *at the bottom,* $e_+$ *and* $e_-$ *as two incomparable middle elements, and* $1$ *at the top.*

**Definition 5.2** (Truth order $\leq_t$ **[Established]**). *The* truth order $\leq_t$ *on* $B_\sigma(\mathbb{D})$ *is the partial order generated by the two covering relations*

$$e_- \leq_t 0, \qquad e_- \leq_t 1, \qquad 0 \leq_t e_+, \qquad 1 \leq_t e_+,$$

*with* $0$ *and* $1$ incomparable *in* $\leq_t$*. The Hasse diagram of* $(B_\sigma(\mathbb{D}), \leq_t)$ *is the four-element diamond with* $e_-$ *at the bottom,* $0$ *and* $1$ *as two incomparable middle elements, and* $e_+$ *at the top.*

Semantically: the information order rises from the *absence of information* (Neither $= 0$: no tail has yet stabilised) through the two unambiguous single-sector answers (True $= e_+$ and False $= e_-$) to the *presence of all information* (Both $= 1$: both sectors simultaneously hold). The truth order runs from *most false* ($e_-$) to *most true* ($e_+$), with $0$ and $1$ as the two *incomparably valid* middle cases (neither and both have the same truth-order status: they are *true-valued as much as false-valued*, differing only in informational content).

**Proposition 5.3** (Lattice structure in each order **[Established]**). *Both* $(B_\sigma(\mathbb{D}), \leq_i)$ *and* $(B_\sigma(\mathbb{D}), \leq_t)$ *are lattices:*

(i) *In $\leq_i$, meet is written $\otimes$ and join is written $\oplus$:*

$$0 = e_+ \otimes e_-, \qquad 1 = e_+ \oplus e_-, \qquad 0 \otimes x = 0, \qquad 1 \oplus x = 1 \quad (x \in B_\sigma(\mathbb{D})).$$

*Call $\otimes$ the* consensus *and $\oplus$ the* pooling *of information.*

(ii) *In $\leq_t$, meet is written $\wedge$ and join is written $\vee$ (the familiar logical conjunction and disjunction):*

$$e_- = 0 \wedge 1, \qquad e_+ = 0 \vee 1, \qquad e_- \wedge x = e_-, \qquad e_+ \vee x = e_+ \quad (x \in B_\sigma(\mathbb{D})).$$

*Each lattice is distributive.*

*Proof.* A four-element diamond with exactly two incomparable middle elements is always a distributive lattice (the only non-distributive four-element lattice is the diamond $M_3$, which has three incomparable middle elements). Routine case-checking verifies the meets and joins stated. The consensus $\otimes$ of two middle elements is the bottom $0$ because $e_+$ and $e_-$ share no positive information; the pooling $\oplus$ of two middle elements is the top $1$ because pooling combines their information. Dually for $\leq_t$, the truth-order meet of the two *incomparable middle elements* $\{0, 1\}$ is the bottom $e_-$, and the truth-order join is the top $e_+$. □

**Theorem 5.4** ($B_\sigma(\mathbb{D})$ is a distributive bilattice **[Established]**). *The quadruple $(B_\sigma(\mathbb{D}), \leq_i, \leq_t)$ is a* distributive bilattice*: both $(B_\sigma(\mathbb{D}), \leq_i)$ and $(B_\sigma(\mathbb{D}), \leq_t)$ are distributive lattices, and the two orders are* compatible *in the sense that each of $\otimes, \oplus, \wedge, \vee$ is monotone with respect to the other order. Explicitly:*

(B1) *$x \leq_i y$ implies $x \wedge z \leq_i y \wedge z$ and $x \vee z \leq_i y \vee z$.*

(B2) *$x \leq_t y$ implies $x \otimes z \leq_t y \otimes z$ and $x \oplus z \leq_t y \oplus z$.*

*In particular the interlaced pair $(\otimes, \oplus, \wedge, \vee)$ satisfies the twelve distributive laws of a Ginsberg–Fitting distributive bilattice.*

*Proof.* Both lattices were shown distributive in Proposition 5.3. Compatibility (B1) and (B2) is verified case-by-case on the sixty-four triples $(x, y, z) \in B_\sigma(\mathbb{D})^3$; because $|B_\sigma(\mathbb{D})| = 4$ the check is finite. We exhibit one illustrative subcase: set $x = 0$, $y = e_+$, $z = e_-$. Then $x \leq_i y$ (from Definition 5.1), and $x \wedge z = 0 \wedge e_- = e_-$ while $y \wedge z = e_+ \wedge e_- = e_-$, so indeed $x \wedge z \leq_i y \wedge z$ (in fact equality). The remaining sixty-three cases are similar and introduce no new phenomena. □

**Remark 5.5** (Bilattice and the $\sigma$-involution)**.** The $\sigma$-involution of Hinge 4 is a bilattice *anti-automorphism* with respect to $\leq_t$ and a bilattice *automorphism* with respect to $\leq_i$. Concretely, $\sigma$ swaps $e_+ \leftrightarrow e_-$ and fixes $0, 1$:

- In $\leq_t$, $\sigma$ is the order-reversing map (the unique non-trivial lattice anti-automorphism of the truth-diamond) and therefore exchanges $\wedge$ and $\vee$: $\sigma(x \wedge y) = \sigma(x) \vee \sigma(y)$.
- In $\leq_i$, $\sigma$ is order-preserving (it swaps two elements at the same height) and commutes with both $\otimes$ and $\oplus$.

Thus $\sigma$ plays the role of *negation* in the truth order while preserving *informational status* — a paraconsistent negation that does not flip informational content.

## 5.2 The four truth values: semantic reading

Each element of $B_\sigma(\mathbb{D})$ receives a canonical linguistic label that tracks its role in the paraconsistent interpretation of propositions. We fix these labels once and for all.

**Definition 5.6** (The four truth values of $\mathsf{Truth}_4$ **[Established]**)**.** *Define* $\mathsf{Truth}_4 := B_\sigma(\mathbb{D}) = \{0, e_+, e_-, 1\}$ *and the canonical labelling*

$$\text{Neither} := 0 \in B_\sigma(\mathbb{D}), \tag{22}$$

$$\text{True} := e_+ \in B_\sigma(\mathbb{D}), \tag{23}$$

$$\text{False} := e_- \in B_\sigma(\mathbb{D}), \tag{24}$$

$$\text{Both} := 1 \in B_\sigma(\mathbb{D}). \tag{25}$$

*We refer to elements of* $\mathsf{Truth}_4$ *interchangeably by their algebraic names* $\{0, e_+, e_-, 1\}$ *and their semantic names* $\{\text{Neither, True, False, Both}\}$.

The semantic readings of the four labels are as follows.

Neither $= 0$**: ontic indeterminacy..** The bottom element $0 \in \mathbb{D}$ is the additive identity. Semantically, $[\![p]\!] =$ Neither records that the Cauchy sequence of tail-prefixes of the $\omega$-germ attached to the proposition $p$ has *not yet stabilised* within the finite-witness budget. This is *ontic*, not *epistemic*: it does not mean "we do not know the truth value of $p$," but rather "there is not yet a truth value for $p$ to know." The stabilisation, when it occurs, produces one of $\{$True, False, Both$\}$; until then, $p$ has no truth-sector home. This is the Panta Rhei resolution of epistemic-versus-ontic uncertainty: epistemic uncertainty asks which classical truth value $p$ has; ontic uncertainty records that $p$ has no classical truth value at all, only an $\omega$-germ that may or may not stabilise.

True $= e_+$**: plus-lobe idempotent..** Semantically, $[\![p]\!] =$ True records that $p$ holds unambiguously and $\neg p$ does not. Algebraically, $e_+$ is the plus-lobe idempotent of $\mathbb{D}$: it satisfies $e_+^2 = e_+$ and $e_+ \cdot e_- = 0$, so a True-valued proposition is $\sigma$-polarised entirely into one sector.

False $= e_-$**: minus-lobe idempotent..** Dually, $[\![p]\!] =$ False records that $\neg p$ holds unambiguously and $p$ does not; algebraically, $e_-$ is the minus-lobe idempotent, and the $\sigma$-involution of Hinge 4 exchanges $e_+ \leftrightarrow e_-$, realising the paraconsistent negation as the lobe-swap on the truth-sector pair.

Both $= 1$**: the idempotent unit..** The top element $1 \in \mathbb{D}$ is the multiplicative identity of the algebra and the sum $e_+ + e_- = 1$ of the two lobe idempotents. Semantically, $[\![p]\!] =$ Both records that both $p$ and $\neg p$ simultaneously hold in the $\omega$-germ stabilisation of a self-referential or paradoxical proposition. This is the paraconsistent fixed point: the Liar proposition $L \leftrightarrow \neg L$ stabilises at $[\![L]\!] =$ Both $= 1$, because the $\omega$-germ of the Liar iteration populates both the plus-lobe and the minus-lobe in equal measure. The algebraic identity Both $= 1$ (the *idempotent unit* of Theorem 1.4) realises the Hegelian "unity of opposites" as an algebraic fact, not a rhetorical figure.

**Remark 5.7** (Information-order reading of the four values)**.** The information order $\leq_i$ orders the four values by how much information they carry: Neither is minimal (no information), True and False each carry one sector's worth, Both carries both sectors' worth. An analogy with the Belnap database interpretation [2]: Neither is an empty record; True and False are one-sided records; Both is a record with conflicting information from two sources. The truth order $\leq_t$, by contrast, orders values by how much they vote for "$p$ holds": False $\leq_t \{0, 1\} \leq_t$ True, with Neither and Both equi-valent in truth-order vote (each votes half-for, half-against, but from opposite epistemic stances).

### 5.3 The paraconsistent connectives

We now define the four connectives that realise the internal logic of $\mathbf{Cat}_\tau$ on the classifier $\Omega_\tau = B_\sigma(\mathbb{D})$.

**Definition 5.8** (Paraconsistent connectives on $\mathsf{Truth}_4$ **[Established]**)**.** *Define the following four operations on* $\mathsf{Truth}_4 = B_\sigma(\mathbb{D})$*:*

(a) Conjunction $\wedge\colon B_\sigma(\mathbb{D}) \times B_\sigma(\mathbb{D}) \to B_\sigma(\mathbb{D})$*: the truth-order meet,*

$$p \wedge q \;:=\; \inf_{\leq_t}\{p, q\}.$$

(b) Disjunction $\vee\colon B_\sigma(\mathbb{D}) \times B_\sigma(\mathbb{D}) \to B_\sigma(\mathbb{D})$*: the truth-order join,*

$$p \vee q \;:=\; \sup_{\leq_t}\{p, q\}.$$

(c) Negation $\neg\colon B_\sigma(\mathbb{D}) \to B_\sigma(\mathbb{D})$*: the $\sigma$-involution of Hinge 4 restricted to* $B_\sigma(\mathbb{D})$*,*

$$\neg p \;:=\; \sigma(p), \qquad \sigma(0) = 0,\ \sigma(e_+) = e_-,\ \sigma(e_-) = e_+,\ \sigma(1) = 1.$$

(d) Implication $\to\colon B_\sigma(\mathbb{D}) \times B_\sigma(\mathbb{D}) \to B_\sigma(\mathbb{D})$*: material-implication extension,*

$$p \to q \;:=\; \neg p \vee q.$$

We record the truth tables explicitly. The proofs are finite case-checks in the distributive bilattice of Theorem 5.4.

**Proposition 5.9** (Truth tables **[Established]**)**.** *The four connectives of Definition 5.8 act on* $B_\sigma(\mathbb{D})$ *as follows.*

(a) Negation $\neg$:

| $p$ | $\neg p$ |
|---|---|
| Neither $= 0$ | Neither $= 0$ |
| True $= e_+$ | False $= e_-$ |
| False $= e_-$ | True $= e_+$ |
| Both $= 1$ | Both $= 1$ |

(b) Conjunction $\wedge$ (truth-order meet):

| $p \wedge q$ | False | Neither | Both | True |
|---|---|---|---|---|
| False | False | False | False | False |
| Neither | False | Neither | False | Neither |
| Both | False | False | Both | Both |
| True | False | Neither | Both | True |

(c) Disjunction $\vee$ (truth-order join):

| $p \vee q$ | False | Neither | Both | True |
|---|---|---|---|---|
| False | False | Neither | Both | True |
| Neither | Neither | Neither | True | True |
| Both | Both | True | Both | True |
| True | True | True | True | True |

(d) Implication $\to$ ($p \to q = \neg p \vee q$):

| $p \to q$ | False | Neither | Both | True |
|---|---|---|---|---|
| False | True | True | True | True |
| Neither | Neither | Neither | True | True |
| Both | Both | True | Both | True |
| True | False | Neither | Both | True |

*Proof.* *(a)* Negation is $\sigma|_{B_\sigma(\mathbb{D})}$ by Hinge 4 [16]: $\sigma$ fixes the trivial idempotents $\{0, 1\}$ and swaps the non-trivial pair $\{e_+, e_-\}$.

*(b,c)* Compute directly from the truth-order lattice $(B_\sigma(\mathbb{D}), \leq_t)$ of Definition 5.2. For example, Both $\wedge$ True: in $\leq_t$ the only lower bound of $\{\text{Both}, \text{True}\}$ that is $\leq_t$ both is Both itself (since Both $\leq_t$ True and Both $\leq_t$ Both), so Both $\wedge$ True $=$ Both. Dually Both $\vee$ True $=$ True. The row/column for Neither arises similarly from the twin relations Neither $\leq_t$ True and False $\leq_t$ Neither (from Definition 5.2).

*(d)* Apply (a) then (c): for example Both $\to$ False $= \neg$ Both $\vee$ False $=$ Both $\vee$ False $=$ Both (using the Both-row of the disjunction table). □

**Remark 5.10** (Classical reduction on $\{\text{True}, \text{False}\}$)**.** Restricting the four connectives to the two-element subset $\{\text{True}, \text{False}\} = \{e_+, e_-\} \subset B_\sigma(\mathbb{D})$ recovers classical two-valued Boolean logic. On $\{\text{True}, \text{False}\}$: $\neg$ exchanges the two; $\wedge$ and $\vee$ are classical conjunction and disjunction; $\to$ is classical material implication. In particular, the classical Boolean sub-logic is the restriction of the $\tau$-topos logic to $\sigma$-non-trivial, non-idempotent-unit propositions. This is the sense in which the $\tau$-topos logic is a *conservative extension* of classical Boolean logic: it agrees with Boolean logic on classical two-valued inputs, and extends it by the two extra values $\{\text{Neither}, \text{Both}\}$ recording ontic indeterminacy and paraconsistent fixed-point stabilisation.

### 5.4 Paraconsistent inference: failure of explosion

We now isolate the key non-classical feature of $\mathsf{Truth}_4$: the failure of *ex contradictione quodlibet* (ECQ), the classical explosion principle $p \wedge \neg p \vdash q$. In a classical Boolean logic ECQ holds vacuously because the premise $p \wedge \neg p = \bot$ is a contradiction from which everything follows; in $\mathsf{Truth}_4$ the premise $p \wedge \neg p$ has a non-trivial truth value at Both, which does not imply arbitrary $q$.

**Proposition 5.11** (Law of Non-Contradiction fails at Both **[Established]**)**.** *In* $\mathsf{Truth}_4$ *the Law of Non-Contradiction* $p \wedge \neg p = \text{Neither}$ *is* not *valid. Specifically,*

$$\text{Both} \wedge \neg\, \text{Both} \;=\; \text{Both} \wedge \text{Both} \;=\; \text{Both} \;\neq\; \text{Neither},$$

*so the classical instance* $p \wedge \neg p \to \bot$ *fails at* $p = \text{Both}$.

*Proof.* By Proposition 5.9(a), $\neg\,\text{Both} = \text{Both}$ (since $\sigma(1) = 1$). By Proposition 5.9(b), the truth-order meet of Both with itself is Both (diagonal entry of the $\wedge$-table), since $\text{Both} \leq_t \text{Both}$ trivially. Hence $\text{Both} \wedge \neg\,\text{Both} = \text{Both}$, which is not the classical Neither (here playing the role of $\bot$). This is the paraconsistent content: a proposition valued Both is a *contradiction* in the classical sense ($p$ and $\neg p$ both hold) but does not collapse to the bottom of the truth-order; it stabilises at the idempotent unit $1$, which is a distinct, positively populated truth value. □

**Proposition 5.12** (Law of Excluded Middle fails at Neither **[Established]**)**.** *In* $\mathsf{Truth}_4$ *the Law of Excluded Middle* $p \vee \neg p = \text{True}$ *is* not *valid. Specifically,*

$$\text{Neither} \vee \neg\, \text{Neither} \;=\; \text{Neither} \vee \text{Neither} \;=\; \text{Neither} \;\neq\; \text{True},$$

*so the classical instance* $p \vee \neg p = \top$ *fails at* $p = \text{Neither}$.

*Proof.* By Proposition 5.9(a), $\neg\,\text{Neither} = \text{Neither}$ (since $\sigma(0) = 0$). By Proposition 5.9(c), the truth-order join of Neither with itself is Neither (diagonal entry of the $\vee$-table). Hence $\text{Neither} \vee \neg\,\text{Neither} = \text{Neither}$, not True. Interpretation: a proposition in the *not-yet-stabilised* state Neither does *not* obey excluded middle, because neither the plus-lobe nor the minus-lobe has been populated by the $\omega$-germ iteration yet. □

**Theorem 5.13** (Paraconsistent non-explosion (bilattice version) **[$\tau$-Effective]**)**.** *In the internal logic of* $\mathbf{Cat}_\tau$*, the classical principle of* ex contradictione quodlibet $p \wedge \neg p \vdash q$ *fails under the FDE designated-preservation entailment of Definition 6.6 (§6). Specifically, taking p with* $[\![p]\!] = \text{Both}$ *gives* $[\![p \wedge \neg p]\!] = \text{Both} \in \mathsf{D} = \{\text{True}, \text{Both}\}$*, and designated- preservation forces* $[\![q]\!] \in \mathsf{D}$ *for the entailment to hold. For q chosen with* $[\![q]\!] = \text{False}$ *or* $[\![q]\!] = \text{Neither}$ *(both outside* $\mathsf{D}$*), the entailment* $\text{Both} \vdash q$ *fails. In particular* Both *does not entail arbitrary q; the classical explosion is blocked.*

*Proof.* We use the FDE designated-preservation entailment of Definition 6.6 (ii): $p \vdash q$ holds iff, at every generalised point $\bar{x}$, $[\![p]\!](\bar{x}) \in \mathsf{D}$ implies $[\![q]\!](\bar{x}) \in \mathsf{D}$, where $\mathsf{D} = \{\text{True}, \text{Both}\} = \{e_+, 1\}$ is the designated truth sector. By Proposition 5.11, $[\![\text{Both} \wedge \neg\,\text{Both}]\!] = \text{Both} \wedge \text{Both} = \text{Both} \in \mathsf{D}$. Hence for $\text{Both} \vdash q$ to hold, we require $[\![q]\!] \in \mathsf{D}$ everywhere. Taking any $q$ with $[\![q]\!] = \text{False} \notin \mathsf{D}$ or $[\![q]\!] = \text{Neither} \notin \mathsf{D}$, the entailment fails: at that $q$, the designated-preservation condition "$\text{Both} \in \mathsf{D} \Rightarrow [\![q]\!] \in \mathsf{D}$" is violated. This is precisely the failure of ECQ: a proposition whose classical contradiction-form $p \wedge \neg p$ is designated does not thereby entail arbitrary $q$. □

**Remark 5.14** (Truth-order reading coincides here)**.** The truth-order entailment $p \vdash_t q := [\![p]\!] \leq_t [\![q]\!]$ is a distinct relation from the FDE designated- preservation entailment of Definition 6.6 (ii) (they differ, e.g., on $\text{Neither} \vdash \text{False}$: $\leq_t$ fails since $\text{False} \leq_t \text{Neither}$, while FDE holds vacuously since $\text{Neither} \notin \mathsf{D}$). However, for the present ECQ computation with $[\![p \wedge \neg p]\!] = \text{Both}$, both readings give the same answer: the $\leq_t$-upper set of Both is $\{\text{Both}, \text{True}\} = \mathsf{D}$, coinciding with the designated sector. We adopt the FDE reading as primary throughout Hinge 6 (following Definition 6.6); the truth-order reading is retained only as an auxiliary tool for bilattice- algebraic verifications where it happens to coincide.

**Remark 5.15** (Why the logic is paraconsistent and not merely three-valued)**.** Three-valued Kleene–Priest logics (such as Priest's $LP$ [27]) add only the single value Both (or Neither, depending on the variant) to the classical $\{\text{True}, \text{False}\}$. The $\tau$-topos logic adds *both* — and critically, it adds them with distinct algebraic identities: Neither $= 0$ is the additive zero while Both $= 1$ is the multiplicative unit. The distinction is structural: Neither $\cdot p = 0$ annihilates any proposition (no tail has stabilised), while Both $\cdot p = p$ is the identity (both sectors' stabilisation contains full information). Hence $\mathsf{Truth}_4$ cannot be collapsed to a three-valued logic without losing this algebraic distinction between *no truth* (0) and *all truth* (1).

### 5.5 Isomorphism with Belnap–Dunn 4

Belnap's original four-valued logic **4** [2], subsequently axiomatised in Dunn's first-degree entailment and treated textbook-style in Priest [27], has truth values $\{\bot, \top, \mathrm{T}, \mathrm{F}\}$ standing respectively for *neither*, *both*, *true*, *false*, with a bilattice structure given by the information and truth orders. We now make the comparison precise.

**Theorem 5.16** (Belnap–Dunn isomorphism **[Established]**)**.** *There is a unique bilattice isomorphism*

$$\Phi \colon \mathbf{4} \longrightarrow \mathsf{Truth}_4 = B_\sigma(\mathbb{D}), \qquad \bot \mapsto 0,\ \top \mapsto 1,\ \mathrm{T} \mapsto e_+,\ \mathrm{F} \mapsto e_-.$$

*Equivalently, the correspondence* $\{\bot, \top, \mathrm{T}, \mathrm{F}\} \leftrightarrow \{\text{Neither}, \text{Both}, \text{True}, \text{False}\}$ *matches Belnap's bilattice structure on the left with the bilattice structure of Theorem 5.4 on the right:*

- $\Phi$ *preserves the information order: Belnap's* $\leq_k$ *("knowledge" order,* $\bot \leq \mathrm{T}, \mathrm{F} \leq \top$*) matches our* $\leq_i$ *pointwise.*
- $\Phi$ *preserves the truth order: Belnap's* $\leq_t$ *(truth order,* $\mathrm{F} \leq \bot, \top \leq \mathrm{T}$*) matches our* $\leq_t$ *pointwise.*
- $\Phi$ *preserves all four connectives: Belnap's* $\neg, \wedge, \vee$ *match our* $\neg, \wedge, \vee$ *under* $\Phi$*, and Belnap's Dunn-style material implication* $\to$ *matches our* $\to$*.*

*The isomorphism is unique because the four truth values are distinguished by their distinct roles in the two orders:* $\bot$ *is the unique* $\leq_i$*-minimum,* $\top$ *is the* $\leq_i$*-maximum,* $\mathrm{T}$ *is the* $\leq_t$*-maximum, and* $\mathrm{F}$ *is the* $\leq_t$*-minimum; each of these properties pins* $\Phi$ *on one generator.*

*Proof.* *Existence.* Define $\Phi$ by the stated assignment on the four generators; it is bijective on underlying sets because both sides have four elements. To see $\Phi$ preserves the two orders, note from Definitions 5.1–5.2 that $(B_\sigma(\mathbb{D}), \leq_i)$ has 0 at the bottom, 1 at the top, and $\{e_+, e_-\}$ as incomparable middles; and $(B_\sigma(\mathbb{D}), \leq_t)$ has $e_-$ at the bottom, $e_+$ at the top, and $\{0, 1\}$ as incomparable middles. Belnap's **4** has the same two orders with the assignments given [2]. The map $\Phi$ matches generators to generators and preserves both orders by inspection of the Hasse diagrams.

Preservation of connectives: $\wedge$ and $\vee$ are the two truth-order lattice operations, which are preserved by any bilattice isomorphism preserving $\leq_t$. Negation $\neg = \sigma|_{B_\sigma(\mathbb{D})}$ matches Belnap's bilattice negation (the unique lattice *anti-automorphism* of $\leq_t$ that is an automorphism of $\leq_i$) because $\sigma$ fixes $\{0, 1\}$ and swaps $\{e_+, e_-\}$ — the same action as Belnap's $\neg$ which fixes $\{\bot, \top\}$ and swaps $\{\mathrm{T}, \mathrm{F}\}$. Material implication $\to$ is then preserved as the composite $\neg p \vee q$.

*Uniqueness.* Any bilattice isomorphism $\Phi$ must send $\leq_i$-minimum to $\leq_i$-minimum, $\leq_i$-maximum to $\leq_i$-maximum, $\leq_t$-minimum to $\leq_t$-minimum, $\leq_t$-maximum to $\leq_t$-maximum. Belnap's $\bot$ is the unique $\leq_i$-minimum and must go to the unique $\leq_i$-minimum of $B_\sigma(\mathbb{D})$, which is $0 =$ Neither. Similarly $\top \mapsto 1$, $\mathrm{T} \mapsto e_+$, $\mathrm{F} \mapsto e_-$; each choice is forced. □

**Remark 5.17** (Significance of the isomorphism)**.** Theorem 5.16 shows that the four-valued internal logic of $\mathbf{Cat}_\tau$ is *not* a novel paraconsistent logic invented for the $\tau$-framework; it is *the same* four-valued logic that Belnap and Dunn introduced fifty years ago for the purposes of database semantics and relevance logic. What is novel in the $\tau$-framework is not the logic itself but its *origin*: the four truth values arise from the split-complex idempotent algebra $\mathbb{D}$ of Hinge 4, not from axiomatic stipulation. The paraconsistent logic of $\mathbf{Cat}_\tau$ is therefore *earned*, not chosen — the same Belnap–Dunn logic, but with a categorical foundation.

### 5.6 Semantic interpretation brackets $[\![-]\!]$

The interpretation of propositions in $\mathbf{Cat}_\tau$ is handled by a semantic-interpretation map $[\![-]\!]$ that sends each proposition to its truth value in $\Omega_\tau = B_\sigma(\mathbb{D})$. We now formalise this map and verify that it commutes with the four connectives in the Tarski style.

**Definition 5.18** (Propositions in $\mathbf{Cat}_\tau$ **[$\tau$-Effective]**)**.** *A* proposition on a carrier $X$ *is a morphism* $p\colon X \to \Omega_\tau$ *in* $\mathbf{Cat}_\tau$*; the set of all propositions on* $X$ *is*

$$\mathrm{Prop}_\tau(X) \;:=\; \mathrm{Hom}_{\mathbf{Cat}_\tau}(X, \Omega_\tau).$$

*The global proposition-type is*

$$\mathrm{Prop}_\tau \;:=\; \mathrm{Prop}_\tau(1) \;=\; \mathrm{Hom}_{\mathbf{Cat}_\tau}(1, \Omega_\tau) \;\cong\; B_\sigma(\mathbb{D}),$$

*where* $1$ *is the terminal object of* $\mathbf{Cat}_\tau$ *from §3.*

For a proposition $p \in \mathrm{Prop}_\tau(X)$ viewed as a morphism $p\colon X \to \Omega_\tau$, the semantic-interpretation bracket $[\![p]\!]$ is defined pointwise on $X$; for a global proposition $p \in \mathrm{Prop}_\tau$, the bracket is just the value of $p$ in $B_\sigma(\mathbb{D})$.

**Definition 5.19** (Semantic-interpretation brackets **[$\tau$-Effective]**)**.** *The* semantic-interpretation brackets *are the map*

$$[\![-]\!]\colon \mathrm{Prop}_\tau \longrightarrow \Omega_\tau \cong B_\sigma(\mathbb{D}), \qquad p \longmapsto [\![p]\!] \in B_\sigma(\mathbb{D}),$$

*assigning to each global proposition its value in the classifier. For propositions on a general carrier* $X$*, the brackets* $[\![-]\!]_X\colon \mathrm{Prop}_\tau(X) \to \Omega_\tau^X$ *are the pointwise extension,* $[\![p]\!]_X(x) := p(x) \in B_\sigma(\mathbb{D})$.

**Proposition 5.20** (Tarski clauses for $[\![-]\!]$ **[$\tau$-Effective]**)**.** *The semantic-interpretation brackets satisfy the standard Tarski-style homomorphism clauses with respect to the four paraconsistent connectives: for all* $p, q \in \mathrm{Prop}_\tau$*,*

$$[\![p \wedge q]\!] \;=\; [\![p]\!] \wedge [\![q]\!], \tag{26}$$
$$[\![p \vee q]\!] \;=\; [\![p]\!] \vee [\![q]\!], \tag{27}$$
$$[\![\neg p]\!] \;=\; \sigma([\![p]\!]) \;=\; \neg[\![p]\!], \tag{28}$$
$$[\![p \to q]\!] \;=\; [\![p]\!] \to [\![q]\!]. \tag{29}$$

*Proof.* Equations (26)–(27) are the definition of the internal conjunction and disjunction as the pointwise truth-order meet and join of the characteristic morphisms via the subobject-classifier structure of §4; see Definition 5.8. Equation (28) is the definition of internal negation as the $\sigma$-involution on the classifier (Definition 5.8(c)). Equation (29) follows from (27) and (28): $[\![p \to q]\!] = [\![\neg p \vee q]\!] = [\![\neg p]\!] \vee [\![q]\!] = \neg[\![p]\!] \vee [\![q]\!] = [\![p]\!] \to [\![q]\!]$. □

**Remark 5.21** (Soundness versus completeness)**.** Proposition 5.20 is the *homomorphism content* of the semantic-interpretation map. The stronger *soundness* statement — that the paraconsistent Belnap–Dunn inference rules are valid under $[\![-]\!]$ — is Theorem 1.2 of §6, which builds on the present section by verifying the paraconsistent inference rules one by one. Completeness (every Belnap–Dunn-valid entailment arises from the internal logic of some object in $\mathbf{Cat}_\tau$) is deferred to Book II [9].

### 5.7 Heyting, Boolean, and paraconsistent internal logic

The four-valued internal logic of $\mathbf{Cat}_\tau$ departs from classical topos theory in a structurally significant way. We close the section by clarifying the relationship.

**Classical topos internal logic (Heyting)..** Lawvere–Tierney's elementary topoi have an internal logic whose truth-object $\Omega$ is a *Heyting algebra*. Intuitionistic propositional connectives — $\wedge$, $\vee$, $\to$, $\neg$, $\bot$, $\top$ — are the Heyting operations, and the internal logic is intuitionistic (constructively-classical): the Law of Excluded Middle need not hold, but the Law of Non-Contradiction always does, and *ex contradictione quodlibet* is valid. The Heyting algebra on $\Omega$ is automatically distributive and complemented in the constructive sense.

**Boolean topoi..** An elementary topos whose $\Omega$ is a *Boolean algebra* is called a *Boolean topos*; its internal logic is classical. Examples: the category of sets, the classifying topos of a Boolean algebra, sheaves on a Stone space.

**$\tau$-topos internal logic (paraconsistent)..** In contrast, **Cat**$_\tau$'s classifier $\Omega_\tau = B_\sigma(\mathbb{D})$ is *neither Heyting nor Boolean* in the classical sense:

- It is not Heyting because LNC fails at Both (Proposition 5.11): there is no Heyting-style pseudo-complement $\neg p$ with $p \wedge \neg p = \bot$ for $p =$ Both.
- It is not Boolean because LEM fails at Neither (Proposition 5.12): there is no complement $\neg p$ with $p \vee \neg p = \top$ for $p =$ Neither.
- It is a distributive bilattice (Theorem 5.4) carrying the paraconsistent Belnap–Dunn structure (Theorem 5.16).

This is a *structural feature*, not a defect. Classical topos theory is the correct framework for mathematics built over Boolean or Heyting foundations. The $\tau$-topos is the correct framework for mathematics built over the split-complex boundary algebra of Hinge 4, which provides a fundamentally different logical primitive — the $\sigma$-equivariant Boolean sublattice $B_\sigma(\mathbb{D})$ — as the algebraic home of truth.

**Theorem 5.22** (**Cat**$_\tau$ is a paraconsistent elementary topos **[$\tau$-Effective]**)**.** ***Cat***$_\tau$ *is an elementary topos in the Lawvere–Tierney sense (finite limits, exponentials, subobject classifier) whose subobject classifier* $\Omega_\tau$ *is a distributive bilattice isomorphic to Belnap–Dunn* **4***, not a Heyting or Boolean algebra. Accordingly its internal logic is a genuine four-valued paraconsistent logic, and* **Cat**$_\tau$ *is a* paraconsistent elementary topos *in the generalised sense of [2, 27]: a Lawvere–Tierney elementary topos whose internal logic is paraconsistent rather than intuitionistic or Boolean.*

*Proof.* The elementary-topos structure is Theorem 1.1; the identification $\Omega_\tau = B_\sigma(\mathbb{D})$ is Definition 5.6 combined with §4. The bilattice structure and the Belnap–Dunn isomorphism are Theorems 5.4 and 5.16. Failure of Heyting / Boolean structure is Propositions 5.11–5.12. These together certify **Cat**$_\tau$ as an elementary topos whose internal logic is paraconsistent in the stated sense. □

**Remark 5.23** (Forward: soundness and circularity)**.** The paraconsistent-soundness theorem (Theorem 1.2) establishes that the Belnap–Dunn inference rules are valid under $[\![-]\!]$; it is proved in §6 using the truth tables and order structure of this section. The circularity-resolution theorem (Theorem 1.3) uses the *fixed-point* properties of {Neither, Both} under $\neg$ — both $\sigma(\text{Neither}) = \text{Neither}$ and $\sigma(\text{Both}) = \text{Both}$ — to stabilise self-referential propositions in §7. The algebraic fact that Both $= 1$ *is the idempotent unit* (Theorem 1.4) is the key structural observation underwriting both soundness and circularity resolution: it tells us that the paraconsistent "both $p$ and $\neg p$" state is an actual algebraic identity, not a failure of classification.

**Remark 5.24** (Registry and Lean preview for §5)**.** The content of this section populates the planned Lean module `TauLib.BookII.Topos.Truth4` (per §1.4) with the following artefacts:

- `Truth4` as a `Fintype`-instance of cardinality 4 with decidable equality, realised as the quotient of `TauLib.BookII.Boundary.Bsigma`.
- `Truth4.leqI`, `Truth4.leqT` as the two `PartialOrder` instances, with the bilattice compatibility (Theorem 5.4) as a `theorem`.
- `Truth4.neg`, `Truth4.meet`, `Truth4.join`, `Truth4.impl` as the four connective operations, with the truth-table correctness lemmas (Proposition 5.9) as `by decide`-verifiable goals.
- `Truth4.lnc_fails`, `Truth4.lem_fails`, `Truth4.non_explosion` (Propositions 5.11–5.12 and Theorem 5.13) as three headline paraconsistency lemmas.
- `Truth4.equiv_Belnap4` as the Belnap isomorphism of Theorem 5.16, realised as an `Equiv` of bilattices.

The algebraic facts (bilattice structure, truth tables, Belnap isomorphism) are all decidable four-element computations (**[Established]**); the semantic facts (Tarski clauses, non-explosion, paraconsistent-topos identification) promote to **[$\tau$-Effective]** via the subobject-classifier semantics of §4.

## 6. PARACONSISTENT SOUNDNESS

### 6.1 Preview

Sections 4–5 constructed the subobject classifier $\Omega_\tau \cong B_\sigma(\mathbb{D}) = \{0, e_+, e_-, 1\}$ and equipped it with the paraconsistent connectives $\wedge, \vee, \neg, \to$ of Belnap–Dunn four-valued logic. The present section verifies that these two constructions mesh: the internal logic of $\mathbf{Cat}_\tau$, read off from the characteristic morphisms $\chi_\tau$, is *sound* with respect to four-valued paraconsistent logic.

We establish Theorem 1.2 in three technical theorems: soundness (Theorem 6.8); failure of explosion (Theorem 6.10); and the classical-subquotient existence (Theorem 4.24). A partial converse at the propositional bilattice level is given in Corollary 6.17; full first-order completeness is deferred to Book II.

The arguments are structural: induction on formula complexity reduces soundness to the bilattice-algebraic clauses, and the $\sigma$-equivariance of $\mathbf{Cat}_\tau$ (inherited from §3 via $\mathbb{D}$) handles negation uniformly.

### 6.2 The internal language of $\mathbf{Cat}_\tau$

**Definition 6.1** (Internal language of $\mathbf{Cat}_\tau$ **[$\tau$-Effective]**). *The* internal language $\mathcal{L}(\mathbf{Cat}_\tau)$ *consists of:*

1. ***Types:*** *objects $X$ of $\mathbf{Cat}_\tau$ (variables written $x : X$).*
2. ***Terms:*** *for each morphism $f\colon X \to Y$, a function symbol $f(-)$ of arity $X \to Y$; terms are built inductively from variables and function symbols.*
3. ***Propositional variables:*** *atomic predicate symbols $p\colon X \to \Omega_\tau$; atomic formulas have the form $p(t)$ with $t$ a term of type $X$.*
4. ***Connectives:*** *the paraconsistent $\wedge, \vee, \neg, \to$ from §5, plus the truth constants $\top, \bot$, True, False, Both, Neither.*
5. ***Quantifiers:*** *$\forall y : Y$ and $\exists y : Y$, interpreted via image-factorisation adjoints to substitution along the projection $X \times^\tau Y \to X$.*

$\mathcal{L}(\mathbf{Cat}_\tau)$ *is a* typed first-order predicate language with four-valued semantics*: the syntactic shape of formulas is classical, but predicates take values in the four-element $\Omega_\tau = B_\sigma(\mathbb{D})$ rather than in the two-element Boolean object.*

**Remark 6.2** (Typed product). The typed product $X \times^\tau Y$ of Definition 3.7 exists whenever $X, Y$ are admissible carriers and has $\sigma$-equivariant projections. Iterated products are written $X^n$, and for free variables $x_i : X_i$ the interpretation ranges over $X_1 \times^\tau \cdots \times^\tau X_n$.

**Remark 6.3** (Closed vs. open formulas). A closed formula $\phi$ has $[\![\phi]\!]\colon 1 \to \Omega_\tau$, i.e. an element of $\Gamma(\Omega_\tau) = B_\sigma(\mathbb{D})$. An open formula $\phi(x_1, \ldots, x_n)$ has $[\![\phi]\!]\colon X_1 \times^\tau \cdots \times^\tau X_n \to \Omega_\tau$.

### 6.3 Semantics: the bracket $[\![-]\!]$

**Definition 6.4** (Semantic interpretation $[\![-]\!]$ **[$\tau$-Effective]**). *For each formula $\phi(x_1, \ldots, x_n)$ of $\mathcal{L}(\mathbf{Cat}_\tau)$ with free variables $x_i : X_i$, the interpretation $[\![\phi]\!]\colon X_1 \times^\tau \cdots \times^\tau X_n \to \Omega_\tau$ is defined by induction on formula complexity:*

1. ***Atoms.*** $[\![p(t)]\!] = p \circ [\![t]\!]$*; for a variable $x$,* $[\![p(x)]\!] = p$.
2. ***Conjunction.*** $[\![\phi \wedge \psi]\!](x) = [\![\phi]\!](x) \wedge_{\Omega_\tau} [\![\psi]\!](x)$.
3. ***Disjunction.*** $[\![\phi \vee \psi]\!](x) = [\![\phi]\!](x) \vee_{\Omega_\tau} [\![\psi]\!](x)$.
4. ***Negation.*** $[\![\neg\phi]\!](x) = \sigma([\![\phi]\!](x))$*, where $\sigma$ is the involution from §2 (swapping $e_+ \leftrightarrow e_-$, fixing $0, 1$).*
5. ***Implication.*** $[\![\phi \to \psi]\!](x) = \sigma([\![\phi]\!](x)) \vee_{\Omega_\tau} [\![\psi]\!](x)$.
6. ***Universal quantifier.*** $[\![\forall y : Y.\, \phi]\!](x) = \bigwedge_{y:Y} [\![\phi]\!](x, y)$*, the fibrewise meet in $\Omega_\tau$ along $\pi_X\colon X \times^\tau Y \to X$.*
7. ***Existential quantifier.*** $[\![\exists y : Y.\, \phi]\!](x) = \bigvee_{y:Y} [\![\phi]\!](x, y)$*, the fibrewise join along the same projection.*
8. ***Truth constants.*** $[\![\text{True}]\!] = e_+$, $[\![\text{False}]\!] = e_-$, $[\![\text{Both}]\!] = 1$, $[\![\text{Neither}]\!] = 0$.

**Remark 6.5** (Welldefinedness and quantifier adjointness). Each clause composes morphisms of $\mathbf{Cat}_\tau$, and $\Omega_\tau$ supports $\wedge, \vee, \sigma, \to$ internally by Definition 5.8 (§5), so $[\![\phi]\!]$ is a well-defined morphism of $\mathbf{Cat}_\tau$. The $\forall/\exists$ clauses exhibit quantification as the right/left adjoints to pullback along $\pi_X$; these operate in the four-valued bilattice rather than in the two-element Boolean algebra, and the required meets/joins exist in $\Omega_\tau$ by the countable boundary-addressed structure of §5.

### 6.4 Validity and entailment

**Definition 6.6** (Designated-value validity and entailment **[$\tau$-Effective]**). *Fix the* designated truth sector $\mathsf{D} := \{\text{True}, \text{Both}\} = \{e_+, 1\} \subset B_\sigma(\mathbb{D})$ *(the Belnap–Dunn "truths and near-truths"). Let* $\phi(x_1, \dots, x_n)$ *be a formula in* $\mathcal{L}(\mathbf{Cat}_\tau)$.

1. $\phi$ *is* valid in $\mathbf{Cat}_\tau$, *written* $\mathbf{Cat}_\tau \vDash \phi$, *iff* $[\![\phi]\!](\bar{x}) \in \mathsf{D}$ *at every generalised point* $\bar{x}$, *i.e.* $[\![\phi]\!](\bar{x}) \in \{e_+, 1\}$. *Equivalently, the characteristic morphism of* $\{\bar{x} : [\![\phi]\!](\bar{x}) \in \mathsf{D}\}$ *coincides with* $\top : 1 \to \Omega_\tau$ *on the whole context.*
2. $\phi$ entails $\psi$, *written* $\phi \vdash \psi$, *iff at every generalised point* $\bar{x}$,

$$[\![\phi]\!](\bar{x}) \in \mathsf{D} \implies [\![\psi]\!](\bar{x}) \in \mathsf{D}.$$

*This is* designated-value preservation *in the sense of Belnap–Dunn first-degree entailment (FDE) [2, 27]:* $\phi \vdash \psi$ *whenever every designated interpretation of* $\phi$ *forces* $\psi$ *to remain designated.*

**Remark 6.7** (Why designated-preservation is the right entailment **[$\tau$-Effective]**). The designated sector $\mathsf{D} = \{\text{True}, \text{Both}\}$ is exactly the set of truth values at which a proposition is "at least classically true modulo paraconsistent fixpoints." The alternative $\{0, e_-\}$-based formulation (where $\phi \vdash \psi$ iff $[\![\phi]\!] \wedge [\![\neg\psi]\!]$ avoids the designated sector) is technically weaker: it permits Both $\vdash$ Neither, because $1 \wedge \sigma(0) = 1 \wedge 0 = 0 \in \{0, e_-\}$, which would make the failure of ECQ less sharp. The FDE designated-preservation formulation of Definition 6.6 (ii) blocks Both $\vdash$ Neither cleanly: if $[\![\phi]\!](\bar{x}) = \text{Both} \in \mathsf{D}$ but $[\![\psi]\!](\bar{x}) = \text{Neither} \notin \mathsf{D}$, the preservation condition fails, so Both $\nvdash$ Neither. This is the sharp paraconsistent-non-explosion that Theorem 6.10 below makes precise.

### 6.5 Main theorem: paraconsistent soundness

**Theorem 6.8** (Paraconsistent soundness of the internal logic of $\mathbf{Cat}_\tau$ **[$\tau$-Effective]**). *The internal logic of* $\mathbf{Cat}_\tau$ *is sound with respect to Belnap–Dunn four-valued propositional logic. That is:*

1. *Every formula derivable in Belnap–Dunn four-valued logic is valid in the* $[\![-]\!]$*-semantics of* $\mathbf{Cat}_\tau$.
2. *Every identity of the bilattice* $B_\sigma(\mathbb{D})$ *(under* $\wedge, \vee, \sigma, 0, 1$*) is internally provable in* $\mathbf{Cat}_\tau$.

*Proof.* Two steps: algebraic verification on $B_\sigma(\mathbb{D})$, then a structural lift.

**Step 1 (bilattice-level Belnap–Dunn axioms).** By Theorem 5.4 (§5), $B_\sigma(\mathbb{D}) = \{0, e_+, e_-, 1\}$ equipped with

- $\wedge$: $0 \wedge x = x \wedge 0 = 0$; $e_+ \wedge e_+ = e_+$; $e_- \wedge e_- = e_-$; $e_+ \wedge e_- = 0$; $1 \wedge x = x$ for $x \in \{e_+, e_-, 1\}$, $1 \wedge 0 = 0$;
- $\vee$: $1 \vee x = 1$; $e_+ \vee e_- = 1$; $0 \vee x = x$;
- $\sigma$: swaps $e_+ \leftrightarrow e_-$, fixes $0, 1$;

is a *De Morgan bilattice* (Ginsberg–Fitting). The Belnap–Dunn axiom list [2] — idempotency, commutativity, associativity, distributivity, $\sigma^2 = \text{id}$, De Morgan ($\sigma(x \wedge y) = \sigma(x) \vee \sigma(y)$ and dually), absorption — reduces to finite tabulations (at most $4^3 = 64$ cases for associativity/distributivity), all of which follow from Proposition 5.9 (§5). This is a finite algebraic identity check, **[Established]**-tier.

**Step 2 (structural induction on formulas).** By induction on formula complexity, Belnap–Dunn-derivable $\phi \Rightarrow \mathbf{Cat}_\tau \vDash \phi$:

- *Atoms.* $\mathbf{Cat}_\tau \vDash p(x)$ iff $p$ factors through $\top$, by Definition 6.6.
- $\wedge$*-introduction/elimination.* If $[\![\phi]\!] = [\![\psi]\!] = e_+$ everywhere, then $[\![\phi \wedge \psi]\!] = e_+ \wedge e_+ = e_+$. The converse uses that $e_+$ is $\wedge$-prime in $B_\sigma(\mathbb{D})$'s truth order: $x \wedge y = e_+$ forces $x = y = e_+$.
- $\vee$*-introduction.* $[\![\phi]\!] = e_+ \Rightarrow [\![\phi \vee \psi]\!] \in \{e_+, 1\}$, which is the designated sector (Remark 6.7).
- $\sigma$*-involutivity.* $\sigma^2 = \text{id}$ gives $[\![\neg\neg\phi]\!] = [\![\phi]\!]$, so double negation holds (unlike in intuitionistic logic).
- *De Morgan and distributivity.* Reduce pointwise to Step 1.
- $\forall$*-introduction.* $\bigwedge_y e_+ = e_+$ by idempotency.
- $\exists$*-introduction.* $[\![\exists y : Y \,.\, \phi]\!] \geq [\![\phi]\!](x, y_0) = e_+$, so lies in $\{e_+, 1\}$, the designated sector.

**Step 3 ($\sigma$-equivariance for negation).** Every $\mathbf{Cat}_\tau$-morphism commutes with $\sigma$ on its source and target when both are $\Omega_\tau$-valued (§3), so $[\![\neg\phi]\!] = \sigma([\![\phi]\!])$ is a morphism of $\mathbf{Cat}_\tau$ (not just of the underlying set). This closes the induction.

Part (ii) specialises part (i) to the free $(\wedge, \vee, \sigma)$-algebra on the four truth constants. □

**Remark 6.9** (Why algebra alone is not enough)**.** Step 1 is a **[Established]**-tier algebraic computation. What promotes it to a soundness theorem for the *internal* logic of $\mathbf{Cat}_\tau$ is Steps 2–3: the categorical semantics $[\![-]\!]$, the $\sigma$-equivariance of morphisms, and the typed product structure. Without the lift, one has only a statement about $B_\sigma(\mathbb{D})$ as an abstract algebra, not about $\mathbf{Cat}_\tau$ as a four-valued topos.

### 6.6 The failure of explosion

**Theorem 6.10** (Failure of explosion in $\mathbf{Cat}_\tau$ **[$\tau$-Effective]**)**.** *In $\mathbf{Cat}_\tau$, the classical rule $p \wedge \neg p \vdash q$ (*ex contradictione quodlibet*) fails under the designated-preservation entailment of Definition 6.6(ii). Specifically, there exist closed propositions $p, q$ with*

$$[\![p \wedge \neg p]\!] = \text{Both} = 1 \ \in \ \mathsf{D}, \qquad [\![q]\!] = \text{Neither} = 0 \ \notin \ \mathsf{D},$$

*and therefore $p \wedge \neg p \nvdash q$. Hence $\mathbf{Cat}_\tau$'s internal logic is genuinely paraconsistent, not classical-with-extra-labels.*

*Proof.* Take $p = L$, the *Liar* of §7 with $L \leftrightarrow \neg L$. By Theorem 1.3, $[\![L]\!] = \text{Both} = 1$ is the $\omega$-germ stabilised value of the template $\Phi(L) = \neg L$. Then

$$[\![L \wedge \neg L]\!] = 1 \wedge_{\Omega_\tau} \sigma(1) = 1 \wedge 1 = 1 = \text{Both},$$

using $\sigma(1) = 1$: both Both and $\neg$ Both evaluate to the idempotent unit, the algebraic form of "unity of opposites." Therefore $[\![L \wedge \neg L]\!] = \text{Both} \in \mathsf{D} = \{e_+, 1\}$.

Take $q$ to be an *unstabilised* propositional symbol, e.g. the fixed point of a template $\Phi$ whose iteration fails to stabilise within the finite-witness budget (§7): by Theorem 7.7 case (d), $[\![q]\!] = \text{Neither} = 0 \notin \mathsf{D}$.

By Definition 6.6(ii), $p \wedge \neg p \ \vdash \ q$ would require that whenever $[\![p \wedge \neg p]\!](\bar{x}) \ \in \ \mathsf{D}$, also $[\![q]\!](\bar{x}) \ \in \ \mathsf{D}$. But here $[\![p \wedge \neg p]\!](\bar{x}) = \text{Both} \in \mathsf{D}$ while $[\![q]\!](\bar{x}) = \text{Neither} \notin \mathsf{D}$. The preservation condition fails, so $p \wedge \neg p \nvdash q$. □

**Remark 6.11** (Why designated-preservation is the sharp formulation)**.** The entailment definition adopted here (Definition 6.6(ii)) is the Belnap–Dunn FDE designated-preservation formulation, which is strictly sharper than the "counterexample set is empty modulo $\{0, e_-\}$" formulation. Under the designated-preservation reading, the non-explosion theorem states directly that $\mathsf{D}$-membership of the premise does not force $\mathsf{D}$-membership of the conclusion — the precise locus where ECQ parts company with the classical rule. The presence of Both in $\Omega_\tau$ does not propagate to every proposition via Belnap–Dunn inference, because the designated-preservation relation respects the four-valued structure of $B_\sigma(\mathbb{D})$.

**Remark 6.12** (Paraconsistency is not incoherence)**.** $\mathbf{Cat}_\tau$ does not contain contradictions; it admits a fourth truth value Both $= 1$ at which a proposition and its negation are jointly supported without the logic then deriving arbitrary $q$. This is the formal content of "paraconsistent."

### 6.7 Classical subquotient on the lobes

Restricted to the truth-order chain $\{e_-, e_+\} = \{\text{False}, \text{True}\}$, the four-valued logic collapses to classical Boolean logic. Strictly, $\{e_+, e_-\}$ is not closed inside $B_\sigma(\mathbb{D})$ under $\wedge$ and $\vee$ ($e_+ \wedge e_- = 0$ and $e_+ \vee e_- = 1$ escape the subset), so the correct "classical collapse" is via a quotient.

**Theorem 6.13** (Classical-topos subquotient of $\mathbf{Cat}_\tau$ **[Established]**)**.** *There is a canonical quotient functor*

$$Q\colon \mathbf{Cat}_\tau \longrightarrow \mathbf{Cat}_\tau^{\text{cl}}$$

*(the classical-sections reflector of §4) such that:*

1. *$\mathbf{Cat}_\tau^{\text{cl}}$ is a classical elementary topos with two-element subobject classifier $\Omega^{\text{cl}} = \{0, 1\}$;*
2. *$Q(\Omega_\tau) \cong \Omega^{\text{cl}}$ is the truth-order truncation sending $\{0, e_-\} \mapsto 0$ and $\{e_+, 1\} \mapsto 1$;*
3. *$\mathbf{Cat}_\tau^{\text{cl}}$'s internal logic is classical (law of excluded middle and ex contradictione quodlibet both hold);*
4. *$Q$ is left adjoint to the inclusion of classical sections; propositions already valued in $\{e_+, e_-\}$ are unchanged by $Q$.*

*Proof.* (1): $Q$ preserves finite limits and exponentials because the truncation factors through the finite-limit-preserving inclusion. (2): Direct computation on the four elements. (3): In $\Omega^{\text{cl}} = \{0, 1\}$, $\neg 0 = 1$ and $\neg 1 = 0$, and $p \vee \neg p = 1$ holds, so both classical rules are derivable; the Belnap–Dunn identities collapse to the classical ones. (4): Universal property of the truncation. ☐

**Remark 6.14** (Subquotient lift). Theorem 4.24 lifts the classifier-level subquotient of §4 from $\Omega_\tau$ (an object) to a full functor of categories. The classical Boolean logic sits inside $\mathbf{Cat}_\tau$'s four-valued logic as the non-paraconsistent fragment.

### 6.8 Partial completeness at the bilattice level

Soundness admits a partial converse at the propositional level; full first-order completeness is deferred.

**Proposition 6.15** (Propositional-level completeness **[τ-Effective]**). *If a formula $\phi$ of $\mathcal{L}(\mathbf{Cat}_\tau)$ built from atomic predicates and $\wedge, \vee, \sigma, \to$ (no quantifiers) satisfies $\mathbf{Cat}_\tau \vDash \phi$, then $\phi$ is derivable in Belnap–Dunn four-valued propositional logic.*

*Sketch.* Belnap–Dunn logic is complete for the class of all De Morgan bilattices (Ginsberg–Fitting), and $B_\sigma(\mathbb{D})$ is a De Morgan bilattice by Step 1 of Theorem 6.8. Hence validity on all $B_\sigma(\mathbb{D})$-valuations (equivalent to $\mathbf{Cat}_\tau \vDash \phi$ at the propositional level by pointwise semantics) implies Belnap–Dunn derivability; see Priest [27] ch. 8 for the underlying completeness theorem. ☐

**Remark 6.16** (First-order completeness deferred). A full first-order completeness theorem for $\mathbf{Cat}_\tau$'s internal logic requires: (a) a Henkin-style construction of canonical models inside $\mathbf{Cat}_\tau$ from consistent theories; (b) a classical-topos embedding preserving modal/quantifier structure; (c) the forthcoming NF confluence theorem of Hinge 7, supplying universal-address resolution for existentially quantified formulas. All three are in flight or planned for Book II [9] via sheaf semantics and Grothendieck-topology refinements.

**Corollary 6.17** (Sound and propositionally complete **[τ-Effective]**). *The propositional fragment of $\mathbf{Cat}_\tau$'s internal logic is sound and complete with respect to Belnap–Dunn four-valued propositional logic: $\mathbf{Cat}_\tau \vDash \phi \iff \vdash_{\mathrm{BD}} \phi$ for every quantifier-free $\phi$.*

*Proof.* Theorem 6.8 (restricted to quantifier-free formulas) paired with Proposition 6.15. ☐

### 6.9 Summary

We have established:

- *Soundness* (Theorem 6.8): Belnap–Dunn-derivable formulas are valid under $[\![-]\!]$.
- *Failure of explosion* (Theorem 6.10): the classical rule $p \wedge \neg p \vdash q$ fails, witnessed by $p = L$ (Liar) and $q$ unstabilised.
- *Classical subquotient* (Theorem 4.24): the two-valued Boolean logic sits inside $\mathbf{Cat}_\tau$'s four-valued logic as the image of the canonical quotient functor $Q$.
- *Propositional completeness* (Corollary 6.17): soundness admits a converse at the quantifier-free level; full first-order completeness is deferred to Book II.

This section realises Theorem 1.2. The next section (§7) uses this soundness foundation to resolve the classical paradoxes of self-reference via constructive $\omega$-germ stabilisation; the existence of $p = L$ with $[\![L]\!] = \mathrm{Both}$ from Theorem 6.10 is the entry point to that resolution.

## 7. CIRCULARITY RESOLUTION VIA $\omega$-GERM STABILISATION

### 7.1 Setup: self-referential propositions as tail-transformers

The classical paradoxes of self-referential semantics — the Liar $L \leftrightarrow \neg L$, Curry's paradox $C \leftrightarrow (C \to \bot)$, the Kleene–Rosser diagonal $\phi \leftrightarrow \neg\mathrm{Prov}(\phi)$ — share a common syntactic shape: a proposition $p$ is identified with the result of applying a propositional template $\Phi\colon \mathrm{Prop}_\tau \to \mathrm{Prop}_\tau$ to $p$ itself. Classical logic has no canonical truth value to assign to such $p$ without either (i) restricting self-reference syntactically (Tarski's meta-language hierarchy) or (ii) adopting an external four-valued axiomatisation (Belnap–Dunn, Priest). In the $\tau$-topos $\mathbf{Cat}_\tau$, however, self-referential fixed points are *computed* — as the $\omega$-germ stabilisation of a Cauchy iteration — and the value they stabilise to lands in the canonical four-element Boolean sublattice $B_\sigma(\mathbb{D}) = \{0, e_+, e_-, 1\}$ (Remark 2.3) internalised by this paper as the subobject classifier $\Omega_\tau$.

We first fix the syntactic class.

**Definition 7.1** (Self-referential proposition **[$\tau$-Effective]**). *A* self-referential proposition *in the internal language of* $\mathbf{Cat}_\tau$ *is a pair* $(p, \Phi)$ *where:*

(a) $\Phi\colon \mathrm{Prop}_\tau \to \mathrm{Prop}_\tau$ *is a* decidable tail-coherent propositional template — *by Remark 2.7, an NF-coded tail transformer* $\Phi \in \mathsf{Hol}_\tau(\Omega_{\mathrm{tail}}, \Omega_{\mathrm{tail}})$ *whose intensional semantics* $[\![\Phi]\!]\colon \Omega_{\mathrm{tail}} \to \Omega_{\mathrm{tail}}$ *is typed,* $\sim$*-stable, and tail-independent at some finite depth* $k_0 \in \mathrm{Idx}$*;*

(b) $p$ *satisfies the fixed-point equation* $p = \Phi(p)$ *up to* $\sim$*-tail equivalence on its NF representative.*

*We write* $\mathrm{SelfRef}_\tau$ *for the class of such pairs and, by a mild abuse of language, refer to "the self-referential proposition* $p = \Phi(p)$*" when the template* $\Phi$ *is clear from context.*

Four canonical templates motivate the class:

$$\Phi_{\mathrm{Liar}}(p) := \neg p, \qquad L = \neg L, \tag{30}$$
$$\Phi_{\mathrm{Curry}}(p) := p \to \bot, \qquad C = (C \to \bot), \tag{31}$$
$$\Phi_{\mathrm{TT}}(p) := p, \qquad T = T, \tag{32}$$
$$\Phi_{\mathrm{KR}}(p) := \neg\,\mathrm{Prov}_\tau(p), \qquad \phi = \neg\,\mathrm{Prov}_\tau(\phi). \tag{33}$$

The negation $\neg$ in (30)–(31) is the paraconsistent Belnap–Dunn negation of §5 (realised on $B_\sigma(\mathbb{D})$ as the involutive swap $e_+ \leftrightarrow e_-, 0 \leftrightarrow 0, 1 \leftrightarrow 1$; concretely $\neg = \sigma$ in the subobject- classifier representation of Remark 2.3). The implication $p \to \bot$ in (31) is the Heyting-implication arrow on $\Omega_\tau$ (§5), which on $B_\sigma(\mathbb{D})$ reduces to the usual Boolean implication $a \to \bot = \neg a$. Each of $\Phi_{\mathrm{Liar}}$, $\Phi_{\mathrm{Curry}}$, $\Phi_{\mathrm{TT}}$, $\Phi_{\mathrm{KR}}$ is a decidable tail-coherent template in the sense of Definition 7.1.

**Remark 7.2** (Why NF-coded templates are the right class **[$\tau$-Effective]**). A syntactic self-reference schema $p \leftrightarrow \Phi(p)$ is meaningful in the $\tau$-topos only when $\Phi$ is itself a transformer admissible in $\mathsf{Hol}_\tau$: otherwise $\Phi(p)$ has no canonical NF representative and the iteration $\Phi^n$ is ill-defined. Definition 7.1 enforces this: the template $\Phi$ must be NF-coded at some finite depth $k_0$ and $\sim$-stable, which are the same admissibility conditions as for any other morphism in $\mathbf{Cat}_\tau$ (Remark 2.7). Classical self-reference schemata that fail admissibility — for instance, those requiring an uncomputable diagonalisation step — simply lie outside $\mathrm{SelfRef}_\tau$ and acquire no truth value in $\Omega_\tau$. This is the first place the $\tau$-kernel discipline intervenes: it filters the class of admissible self-reference *before* the stabilisation procedure begins.

### 7.2 The Cauchy iteration and the $\omega$-germ stabilised value

Each decidable tail-coherent template $\Phi \in \mathsf{Hol}_\tau(\Omega_{\mathrm{tail}}, \Omega_{\mathrm{tail}})$ generates a canonical Cauchy sequence of NF-coded tail transformers.

**Definition 7.3** (Cauchy iteration **[$\tau$-Effective]**). *Fix a* designated iteration seed

$$s_\star \in \Omega_{\mathrm{tail}} \quad \textit{with} \quad [\![s_\star]\!] = \mathrm{False} = e_- \in B_\sigma(\mathbb{D}),$$

*i.e., an NF-coded tail transformer whose tail-stabilised truth value lies in the minus-lobe "classical False" sector (not at the ontic zero* Neither $= 0$*, which would be* $\sigma$*-fixed and yield a trivial iteration). For* $\Phi \in \mathsf{Hol}_\tau(\Omega_{\mathrm{tail}}, \Omega_{\mathrm{tail}})$ *the* Cauchy iteration from seed $s_\star$ *is the sequence* $(\Phi^n(s_\star))_{n\in\mathbb{N}}$ *in* $\mathsf{Hol}_\tau(\Omega_{\mathrm{tail}}, \Omega_{\mathrm{tail}})$ *defined by*

$$\Phi^0(s_\star) := s_\star, \qquad \Phi^{n+1}(s_\star) := \Phi(\Phi^n(s_\star)) \in \Omega_{\mathrm{tail}}. \tag{34}$$

*Under the subobject-classifier semantics* $[\![-]\!]$ *of Definition 5.19 (§5), each* $\Phi^n(s_\star)$ *has a truth value* $[\![\Phi^n(s_\star)]\!] \in B_\sigma(\mathbb{D})$ *computed by the four-valued internal logic of §5. The choice of seed is a calibration convention:* $s_\star = e_-$ *(the classical-False polarity) is canonical; the* $\sigma$*-dual choice* $s_\star = e_+$ *yields the same classification up to the canonical* $\sigma$*-swap on* $B_\sigma(\mathbb{D})$*.*

**Remark 7.4** (Seed vs. ontic zero). The seed $s_\star = e_-$ must not be confused with the ontic zero $0 = $ Neither of $B_\sigma(\mathbb{D})$. The zero element $0 \in \mathbb{D}$ is the additive identity and semantically represents *ontic pre-stabilisation* (no tail-coherence certificate yet); it is $\sigma$-fixed and hence a fixed point of every $\sigma$-equivariant $\Phi$, which would make the iteration constant at Neither for any $\Phi$. The classical-False seed $e_-$, by contrast, is $\sigma$-swapped to $e_+$ by one application of $\neg$, which is what drives the period-2 behaviour of the Liar (Theorem 7.10). This distinction is symbol-critical: we henceforth reserve $s_\star$ for the iteration seed and $0 \in B_\sigma(\mathbb{D})$ for the ontic-zero truth value.

**Definition 7.5** ($\omega$-germ stabilised value **[$\tau$-Effective]**). *For a self-referential proposition* $(p, \Phi) \in \mathrm{SelfRef}_\tau$*, the* $\omega$-germ stabilised truth value *of* $p$ *is*

$$[\![p]\!] \;:=\; \lim_{n\to\infty} [\![\Phi^n(s_\star)]\!] \;\in\; B_\sigma(\mathbb{D}), \tag{35}$$

*where the limit is taken in the* $\sim$*-tail topology on* $\Omega_{\text{tail}}$ *and pushed forward to* $B_\sigma(\mathbb{D})$ *under the sector decomposition* $f \mapsto (f_+, f_-)$ *of Remark 2.12. Equivalently,* $[\![p]\!]$ *is the image in* $B_\sigma(\mathbb{D})$ *of the* $\sim$*-tail equivalence class of the iteration* (34) *under the canonical map* $\Omega_{\text{tail}}/\sim \to B_\sigma(\mathbb{D})$ *induced by the lobe projections* $\pi_\pm : \mathbb{D} \to \mathcal{R}'_\partial$ *(Remark 2.2).*

**Remark 7.6** (Existence of the limit **[$\tau$-Effective]**, modulo Lemma 2.13). The limit (35) is well defined because each $\Phi^n(s_\star)$ is NF-coded and tail-coherent (Definition 7.3), and the $\sim$-tail equivalence classes of the iteration form a Cauchy sequence in the profinite $\sim$-topology on $\Omega_{\text{tail}}/\sim$. The pass-finite convergence follows from the tail-independence of $\Phi$ at depth $k_0 \in \mathrm{Idx}$: after at most $k_0$ iterations the NF representative of $\Phi^n(s_\star)$ has frozen outside its dependency window, and the subsequent sector component $\pi_\pm \circ [\![\Phi^n(s_\star)]\!]$ is either (a) eventually constant in a single idempotent lobe, (b) eventually periodic with period 2, or (c) persistently non-stabilised within any finite witness budget; see Theorem 7.7 below. The identification of (35) with a specific element of $B_\sigma(\mathbb{D})$ depends on Lemma 2.13 to guarantee that the limit is path-independent under NF reduction. Weak confluence at the Hinge-5 level already suffices for the limit to exist as an $\sim$-tail equivalence class; strict confluence, pending from Hinge 7, upgrades the class to a canonical $B_\sigma(\mathbb{D})$-valued truth value. All four-sector classifications below are therefore **[$\tau$-Effective]** *modulo Hinge 7*.

### 7.3 The four-sector classification theorem

The following theorem is the paper's central constructive identification. It states that every self-referential proposition admits a *definite* truth value in $\Omega_\tau = B_\sigma(\mathbb{D})$, computed by the Cauchy iteration (34). The four-way classification is forced by the four-atom structure of $B_\sigma(\mathbb{D})$ itself: the two single-lobe landings, the period-2 oscillation summing to the idempotent unit, and the non-stabilisation that lands on the zero element.

**Theorem 7.7** (Four-sector classification **[$\tau$-Effective]**, modulo Hinge 7 canonical-address NF confluence). *For every self-referential proposition* $(p, \Phi) \in \mathrm{SelfRef}_\tau$*, the* $\omega$*-germ stabilised truth value* $[\![p]\!] \in B_\sigma(\mathbb{D})$ *exists and belongs to one of four mutually exclusive sectors determined by the asymptotic behaviour of the Cauchy iteration* $([\![\Phi^n(s_\star)]\!])_{n\in\mathbb{N}}$ *in* $B_\sigma(\mathbb{D})$*:*

(a) (Single-lobe $e_+$-stabilisation.) *If there exists* $n_0 \in \mathbb{N}$ *such that* $[\![\Phi^n(s_\star)]\!] = e_+$ *for all* $n \geq n_0$*, then*

$$[\![p]\!] \;=\; e_+ \;=\; \mathrm{True}\,.$$

(b) (Single-lobe $e_-$-stabilisation.) *If there exists* $n_0 \in \mathbb{N}$ *such that* $[\![\Phi^n(s_\star)]\!] = e_-$ *for all* $n \geq n_0$*, then*

$$[\![p]\!] \;=\; e_- \;=\; \mathrm{False}\,.$$

(c) (Period-2 oscillation.) *If there exists* $n_0 \in \mathbb{N}$ *such that* $[\![\Phi^{n+1}(s_\star)]\!] \neq [\![\Phi^n(s_\star)]\!]$ *and* $[\![\Phi^{n+2}(s_\star)]\!] = [\![\Phi^n(s_\star)]\!] \in \{e_+, e_-\}$ *for all* $n \geq n_0$*, then the pooled information content of the alternating orbit lands at the information-order join of the two lobes:*

$$[\![p]\!] \;=\; e_+ \oplus e_- \;=\; 1 \;=\; \mathrm{Both},$$

*where* $\oplus$ *is the information-order join on* $B_\sigma(\mathbb{D})$ *of §5 (Proposition 5.9). The idempotent identity* $e_+ + e_- = 1$ *in* $\mathbb{D}$ *coincides with* $e_+ \oplus e_- = 1$ *in the bilattice* $B_\sigma(\mathbb{D})$ *(the two binary operations agree on pairs of orthogonal idempotents), so the sector-decomposition reading and the information-order reading yield the same value.*

(d) (Non-stabilisation.) *If the sequence* $([\![\Phi^n(s_\star)]\!])_{n\in\mathbb{N}}$ *fails to satisfy any of the above three conditions within any finite witness budget bounded by* $k_0$*, then*

$$[\![p]\!] \;=\; 0 \;=\; \mathrm{Neither}\,.$$

*The four cases are exhaustive: the asymptotic behaviour of the iteration falls into exactly one.*

*Proof.* The proof proceeds by reducing each of the four asymptotic behaviours to a specific algebraic identity in the idempotent calculus of $B_\sigma(\mathbb{D})$ and then verifying, via the idempotent- supported holomorphy decomposition of Remark 2.12, that this identity is preserved under the pushforward $\Omega_{\text{tail}}/\sim \to B_\sigma(\mathbb{D})$.

*Cases (a) and (b): single-lobe stabilisation.* Suppose $[\![\Phi^n(s_\star)]\!] = e_+$ eventually. By tail-independence of $\Phi$ at depth $k_0$, the NF representatives of $\Phi^n(s_\star)$ for $n \geq k_0$ agree on their sector component $\pi_+([\![\Phi^n(s_\star)]\!]) = 1_{\mathcal{R}'_\partial}$ and $\pi_-([\![\Phi^n(s_\star)]\!]) = 0_{\mathcal{R}'_\partial}$. The limit of a constant sequence in the $\sim$-tail topology is the constant itself, so

$$[\![p]\!] \;=\; \lim_{n\to\infty} [\![\Phi^n(s_\star)]\!] \;=\; e_+ \cdot 1_{\mathcal{R}'_\partial} + e_- \cdot 0_{\mathcal{R}'_\partial} \;=\; e_+ \;=\; \text{True}\,.$$

Case (b) is symmetric under the $\sigma$-involution of Remark 2.3: applying $\sigma$ to the entire iteration exchanges $e_+ \leftrightarrow e_-$, so eventual $e_-$-stabilisation yields $[\![p]\!] = e_- = \text{False}$ by the same argument.

*Case (c): period-2 oscillation.* Suppose the iteration eventually alternates between $e_+$ and $e_-$ with period 2. In the $\sim$-tail topology, a period-2 oscillation is not convergent at the level of the sequence itself. The $\omega$-germ stabilisation is defined instead as the information-order join of the orbit's pointwise sector contributions. Because the alternating orbit visits both lobe sectors, its pointwise tail-coherence certificate pools information from $e_+$ and $e_-$; the pooled content is the information-order join

$$e_+ \oplus e_- \;=\; 1 \;=\; \text{Both} \;\in\; B_\sigma(\mathbb{D}),$$

where $\oplus$ is the information-order join established in §5 (Proposition 5.9). For the specific case of two orthogonal idempotents, this information-order join coincides with the ring-theoretic sum $e_+ + e_-$ in $\mathbb{D}$ (and with the truth-order join $e_+ \vee e_-$ in the Boolean sublattice $B_\sigma(\mathbb{D})$), a happy coincidence that holds because $B_\sigma(\mathbb{D})$ is a distributive bilattice in which both orders agree on orthogonal-idempotent pairs. Equivalently, the $\omega$-germ stabilisation is not the limit of the sequence but the $\sim$-tail equivalence class of the *joint trajectory*, which contains both $e_+$ and $e_-$ as limit points. The canonical representative of a $\sim$-tail class containing both lobe idempotents in the image of $\pi_+$ and $\pi_-$ is precisely the idempotent unit $e_+ + e_- = 1 \in \mathbb{D}$, because the join of the two single-lobe atoms in the Boolean sublattice $B_\sigma(\mathbb{D})$ is 1. Hence $[\![p]\!] = 1 = \text{Both}$.

The algebraic content of this case is the central identification (36) below: the period-2 oscillation *is* the idempotent unit of $\mathbb{D}$. This is Theorem 1.4 of the introduction, realised here as a concrete output of the Cauchy iteration.

*Case (d): non-stabilisation.* Suppose the iteration fails to enter any of the three stabilised regimes within any finite witness budget bounded by $k_0$. Then the $\sim$-tail equivalence class of the iteration has no canonical NF representative among the three nontrivial atoms $\{e_+, e_-, 1\}$ of $B_\sigma(\mathbb{D})$. By the four-atom completeness of $B_\sigma(\mathbb{D})$ (Remark 2.3), the only remaining atom is $0 = \text{Neither}$, which is the zero element of $\mathbb{D}$ — the formal NF code for "no tail-coherence certificate within the witness budget." This is the ontic realisation of the non- stabilisation: no Cauchy witness, no sector support, no idempotent lobe, hence $[\![p]\!] = 0 = \text{Neither}$.

*Exhaustiveness.* The four cases are mutually exclusive by construction, and they exhaust the possible asymptotic behaviours of a sequence in a finite set $B_\sigma(\mathbb{D}) = \{0, e_+, e_-, 1\}$: either the sequence is eventually constant at one of the three nonzero atoms (cases (a), (b), or (c′) a 1-stabilised sequence which we group with (c) below), or it eventually oscillates with period 2 between two lobes (case (c)), or it fails to stabilise (case (d)). The case of a constant 1-value in the sequence is subsumed by case (c) as a degenerate periodic orbit ($1 = e_+ + e_-$ is the joint limit point), so the four-way partition is complete. □

**Remark 7.8** (The four sectors are algebraically forced **[$\tau$-Effective]**)**.** The four cases of Theorem 7.7 are not a *choice*: they are the four atoms of $B_\sigma(\mathbb{D})$ itself, and the Cauchy iteration acts on $B_\sigma(\mathbb{D})$ by a depth-0 idempotent dynamical system whose orbits are exactly the four classified patterns. Concretely, writing the iteration as an $B_\sigma(\mathbb{D})$-valued map $\Psi\colon B_\sigma(\mathbb{D}) \to B_\sigma(\mathbb{D})$, $\Psi(b) := [\![\Phi(b)]\!]$ (extended from the designated seed $s_\star$ to all of $B_\sigma(\mathbb{D})$), the eventual behaviour of $\Psi^n(s_\star)$ on the four-element state space is completely classified by the theory of finite-state dynamical systems: single-point attractors (cases (a), (b)), 2-cycles (case (c)), and transient orbits that exit the finite-witness budget without stabilising (case (d)). *The four-valued internal logic* $\mathsf{Truth}_4 = B_\sigma(\mathbb{D})$ *is therefore the minimal target space in which every self-referential proposition has a definite truth value.* A three-valued logic (e.g. Kleene's) has no fixed point for the period-2 Liar orbit; a two-valued logic has no room for Neither or Both; a five-or-more-valued logic over-classifies.

**Remark 7.9** (The central algebraic identity $\text{Both} = 1$ **[$\tau$-Effective]**)**.** The step in case (c) of the proof where $e_+ + e_- = 1 = \text{Both}$ is the paper's key algebraic identity:

$$\text{Both} \;=\; e_+ + e_- \;=\; 1 \;\in\; \mathbb{D}. \tag{36}$$

This is *not* an axiomatic stipulation. It is the identity of the idempotent unit of the split-complex algebra $\mathbb{D}$ (Remark 2.2), inherited from Hinge 4 [16], now realised as the output of a period-2 Cauchy iteration on the subobject classifier of $\mathbf{Cat}_\tau$. Theorem 1.4 of §1.3 is precisely this identification, promoted to a main result of the paper. The Hegelian "unity of opposites" — the thesis that $p$ and its negation $\neg p$ are one at the paraconsistent fixed point — realises as the algebraic identity that the sum of the two lobe idempotents is the multiplicative unit of the boundary algebra.

### 7.4 Explicit stabilisation for the canonical examples

We apply Theorem 7.7 to the four canonical templates (30)–(33).

*The Liar: $L = \neg L$ stabilises at* Both

**Theorem 7.10** (Liar lands on Both **[$\tau$-Effective]**, modulo Hinge 7)**.** *For the Liar template $\Phi_{\text{Liar}}(p) = \neg p$, the Cauchy iteration from seed $s_\star = e_-$ is a period-2 oscillation between $e_+$ and $e_-$, and the $\omega$-germ stabilised truth value is*

$$[\![L]\!] \;=\; \text{info-join}_{n\in\mathbb{N}} [\![\Phi^n_{\text{Liar}}(s_\star)]\!] \;=\; e_+ \oplus e_- \;=\; 1 \;=\; \text{Both}, \tag{37}$$

*where $\oplus$ denotes the* information-order join *on $B_\sigma(\mathbb{D})$ (§5, Proposition 5.9): $e_+ \oplus e_- = 1$ because the two lobe sectors, pooled together, are maximally informative. The Liar proposition is genuinely "both true and false" in $\mathbf{Cat}_\tau$'s internal logic: its $\omega$-germ stabilisation is the idempotent unit of the boundary algebra $\mathbb{D}$.*

*Proof.* Under the paraconsistent Belnap–Dunn negation on $B_\sigma(\mathbb{D})$ (realised as the $\sigma$-swap $e_+ \leftrightarrow e_-$, $0 \mapsto 0$, $1 \mapsto 1$; see §5), starting from $[\![s_\star]\!] = e_-$, the iteration unfolds as:

| $n$ | $[\![\Phi^n_{\text{Liar}}(s_\star)]\!]$ |
|---|---|
| 0 | $e_- = \text{False}$ |
| 1 | $\neg e_- = e_+ = \text{True}$ |
| 2 | $\neg e_+ = e_- = \text{False}$ |
| 3 | $\neg e_- = e_+ = \text{True}$ |
| $\vdots$ | $\vdots$ |

The alternating orbit $\{e_+, e_-\}$ has no limit in the truth-order sense (the two values are truth-incomparable), nor any ring-theoretic limit (the sequence oscillates and does not converge in the metric sense). But in the *information order* on $B_\sigma(\mathbb{D})$, the orbit's pointwise information content accumulates: each application of $\neg$ adds the dual lobe's information to the pooled content. The canonical ω-germ limit of a period-2 orbit $\{e_+, e_-\}$ under pooled-information accumulation is therefore the information-order join $e_+ \oplus e_-$, which by Proposition 5.9 equals $1 = \text{Both}$ in $B_\sigma(\mathbb{D})$.

This is case (c) of Theorem 7.7, and the calibration-independence follows from $\sigma$-symmetry: the alternative seed $s_\star = e_+$ gives the orbit $\{e_+, e_-, e_+, \ldots\}$ with the same information-order join, hence the same limit $\text{Both} = 1$. □

**Remark 7.11** (The Liar is the canonical paraconsistent fixed point **[$\tau$-Effective]**)**.** Theorem 7.10 resolves the Liar paradox ontically: $L$ is neither a classically-true proposition nor a classically-false one. It is a proposition whose $\omega$-germ stabilisation equals the idempotent unit of the boundary algebra — the Boolean-lattice top element of $B_\sigma(\mathbb{D})$. Classical logic has no place for such a value (and hence declares $L$ paradoxical); the $\tau$-topos has exactly four truth values, one of which is $\text{Both} = 1$, and the Liar lands there by direct computation.

*Curry's paradox: $C = (C \to \bot)$ and its classification*

Curry's paradox is the schema $C \leftrightarrow (C \to \bot)$: the proposition $C$ is identified with its own negation (since $p \to \bot \equiv \neg p$ in any logic with implication-as- inhabitation). Classically, Curry's paradox permits a derivation of any proposition $q$ from the schema itself, via repeated application of modus ponens; the derivation is the classical "ex falso quodlibet" in its most devastating form, because it does not require an explicit contradiction as a premise. The paraconsistent resolution of §6 blocks this derivation by failing *ex contradictione quodlibet*. What happens in the $\omega$-germ stabilisation?

**Theorem 7.12** (Curry's paradox: $[\![C]\!]$ is Neither or Both, depending on tail-coherence **[$\tau$-Effective]**, modulo Hinge 7)**.** *For the Curry template $\Phi_{\text{Curry}}(p) = p \to \bot$, the Cauchy iteration simplifies to the Liar iteration on $B_\sigma(\mathbb{D})$: $\Phi_{\text{Curry}}$*

*agrees with $\Phi_{\text{Liar}}$ pointwise on $B_\sigma(\mathbb{D})$ (because $b \to \bot = \neg b$ for every $b \in B_\sigma(\mathbb{D})$ under the Heyting-implication structure of §5). Hence:*

(a) If the Curry template is admissible as an NF-coded tail transformer at depth $k_0$ *(i.e. Definition 7.1 applies and the iteration enters the Liar-equivalent period-2 orbit within the witness budget), then* $[\![C]\!] = \text{Both}$ *by Theorem 7.10.*

(b) If the Curry template is *not* admissible as an NF-coded tail transformer — for instance, if the classical Curry derivation is attempted but fails tail-coherence because the Heyting implication $p \to \bot$ cannot be given a uniform finite-witness NF representation on the iterated argument — *then the iteration has no $B_\sigma(\mathbb{D})$-representable limit within the witness budget* $k_0$*, and* $[\![C]\!] = 0 = \text{Neither}$ *by case (d) of Theorem 7.7.*

*The precise classification — (a) or (b) — is determined by the decidability of the NF code for the Curry template. In particular,* the classical Curry derivation of an arbitrary proposition $q$ from $C$ fails in $\mathbf{Cat}_\tau$ *because the paraconsistent soundness theorem (§6) blocks modus ponens from the candidate fixed point* $[\![C]\!]$.

*Proof. Reduction to the Liar.* Throughout this proof, $\bot$ denotes the *classical-False* propositional constant of $\mathbf{Cat}_\tau$, i.e. the unique global section $\bot\colon 1 \to \Omega_\tau$ with $[\![\bot]\!] = \text{False} = e_-$ (not the algebraic zero $0 = \text{Neither}$ of $B_\sigma(\mathbb{D})$). The identity $b \to \bot \equiv \neg b$ then holds for every $b \in \{0, e_+, e_-, 1\}$. Indeed, by the implication table in Proposition 5.9(d) with $q = \text{False}$ (column False):

$$\text{False} \to \text{False} = \text{True}, \quad \text{Neither} \to \text{False} = \text{Neither}, \quad \text{Both} \to \text{False} = \text{Both}, \quad \text{True} \to \text{False} = \text{False},$$

which matches, respectively, $\neg\,\text{False} = \text{True}$, $\neg\,\text{Neither} = \text{Neither}$, $\neg\,\text{Both} = \text{Both}$, and $\neg\,\text{True} = \text{False}$ from Proposition 5.9(a). Therefore $\Phi_{\text{Curry}}(b) = \Phi_{\text{Liar}}(b)$ for every $b \in B_\sigma(\mathbb{D})$, and the Cauchy iteration of $\Phi_{\text{Curry}}$ starting from the designated seed $s_\star$ (with $[\![s_\star]\!] = e_-$ as in Definition 7.3) coincides with the Liar iteration once the Boolean-lattice equality has been internalised.

*Case (a): admissible NF coding.* If $\Phi_{\text{Curry}}$ admits an NF-coded tail-transformer representation at depth $k_0$ — i.e. if the Heyting-implication arrow on $\Omega_\tau$ is implementable as a tail-transformer code on $\Omega_{\text{tail}}$ with finite witness — then the iteration proceeds on $B_\sigma(\mathbb{D})$ as the Liar does, and case (c) of Theorem 7.7 yields $[\![C]\!] = \text{Both}$. (This is the "Curry behaves like Liar" scenario.)

*Case (b): Curry template fails NF admissibility.* The classical Curry derivation — "assume $C$, derive $\bot$ by modus ponens, conclude $\neg C$, then re-assume $C$, ..." — requires the Heyting implication to be applied to a self- referential argument in a way that does not commute with the $\sim$-tail equivalence on $\Omega_{\text{tail}}$. If the implementation of $\Phi_{\text{Curry}}$ fails $\sim$-stability on the iterated argument (for instance, because the Curry derivation escapes the witness budget $k_0$), then the iteration lacks a tail-coherence certificate and lands on $0 = \text{Neither}$ by case (d) of Theorem 7.7. This is the $\mathbf{Cat}_\tau$-analogue of the classical Curry derivation "failing to land anywhere": the paraconsistent filter blocks the explosion, and the $\omega$-germ stabilisation simply declares no truth value.

*Classical Curry derivation blocked.* The classical derivation of an arbitrary $q$ from $C = (C \to \bot)$ uses modus ponens on the premise "$C \wedge (C \to \bot)$" to conclude $\bot$, then ex falso quodlibet to conclude $q$. In $\mathbf{Cat}_\tau$'s internal logic, modus ponens is valid as an inference rule on $\Omega_\tau$, but the step "$\bot \vdash q$" is blocked by FDE designated-preservation (Definition 6.6 (ii), §6): the value $[\![\bot]\!] = e_- = \text{False}$ is *not* designated ($e_- \notin \mathsf{D} = \{e_+, 1\}$), so from "$[\![\bot]\!] \in \mathsf{D}$" being false, nothing about $[\![q]\!]$ is forced. The derivation therefore terminates at $[\![C]\!] \in \{\text{Both}, \text{Neither}\}$ without escaping into the full Boolean lattice of classical truth values. See §6 for the detailed soundness argument. □

**Remark 7.13** (Two resolutions of Curry are possible in the $\tau$-topos **[$\tau$-Effective]**)**.** The branching in Theorem 7.12 is not a defect of the theory. It reflects the fact that Curry's paradox has *two* distinct reading registers:

(i) As a *semantic* paradox (Curry-as-Liar): the template $p \to \bot$ reduces to $\neg p$ and the Liar analysis applies. $[\![C]\!] = \text{Both}$.

(ii) As a *syntactic* paradox (Curry-as-non-normalising $\lambda$-term): the Kleene–Rosser diagonal form of Curry encodes an uncomputable fixed point that fails NF-coding at every finite depth. $[\![C]\!] = \text{Neither}$.

The classical conflation of (i) and (ii) is the root of the classical paradox's devastating character: the *syntactic* non-normalisation is interpreted as a *semantic* fact about truth, and modus ponens is used to derive nonsense. The $\tau$-topos separates the two: (i) yields a paraconsistent fixed point at Both, (ii) yields an ontic non-stabilisation at Neither, and neither of them licences the classical explosion [27].

*The Truth-teller: $T = T$ and its ambiguity*

**Theorem 7.14** (Truth-teller lands on Neither or an ambiguous lobe value **[$\tau$-Effective]**, modulo Hinge 7)**.** *For the identity template $\Phi_{\mathrm{TT}}(p) = p$, the Cauchy iteration is constant: $\Phi^n_{\mathrm{TT}}(s_\star) = s_\star$ for all $n \in \mathbb{N}$. With the designated iteration seed $s_\star = e_-$ of Definition 7.3, the iteration stabilises at $e_-$ and Theorem 7.7(b) yields $[\![T]\!] = \mathrm{False} = e_-$.*

*However, the Truth-teller's fixed-point equation $T = T$ does not* uniquely *specify the initial value of the iteration: any $b_0 \in B_\sigma(\mathbb{D})$ satisfies $\Phi_{\mathrm{TT}}(b_0) = b_0$.* Generalising *the iteration-seed convention of Definition 7.3 to admit an arbitrary $B_\sigma(\mathbb{D})$-valued initial condition $b_0$, the Truth-teller admits* four *$\omega$-germ stabilised truth values: $[\![T]\!] \in \{\mathrm{Neither}, \mathrm{True}, \mathrm{False}, \mathrm{Both}\}$, one per initial condition. The Truth-teller is therefore* ambiguous *in $\mathbf{Cat}_\tau$'s internal logic: its truth value is lobe-choice- dependent, and the paraconsistent four-valued logic reveals this ambiguity as a structural property, not as a paradox.*

*Proof.* The iteration is trivially constant, so Theorem 7.7 applies directly to whichever initial value $b_0$ is chosen. Since $\Phi_{\mathrm{TT}}$ fixes every $b_0 \in B_\sigma(\mathbb{D})$, all four atoms are legitimate stabilised values.

The ambiguity is not a failure of Definition 7.1: it is a genuine *lack of determinacy* in the Truth-teller's semantic content. Classical logic rules the Truth-teller "undefined" or "arbitrary"; the paraconsistent $\tau$-topos records the arbitrariness as a structural fact: any atom of $B_\sigma(\mathbb{D})$ is a fixed point of the Truth-teller template, and no further information disambiguates among them. □

**Remark 7.15** (Truth-teller vs. Liar: the symmetry broken by negation **[$\tau$-Effective]**)**.** The Truth-teller and the Liar are symmetric as fixed-point equations ($T = T$ vs. $L = \neg L$), but their $\omega$-germ stabilisations differ in a crucial way: the Liar's template $\neg$ is *non-trivial* on $B_\sigma(\mathbb{D})$ (it swaps $e_+ \leftrightarrow e_-$), so the iteration is forced into a period-2 orbit and lands on Both unambiguously; the Truth- teller's template is the identity, so the iteration is trivial and lands on whatever initial condition is chosen. The negation in the Liar template therefore plays the role of a "paraconsistent forcing": it compels the iteration into the idempotent unit of the Boolean lattice. Without such forcing, the Truth-teller retains all four atoms as legitimate fixed points.

*Kleene–Rosser: $\phi = \neg\mathrm{Prov}_\tau(\phi)$*

**Theorem 7.16** (Kleene–Rosser diagonal in $\mathbf{Cat}_\tau$ **[$\tau$-Effective]**, modulo Hinge 7)**.** *For the Kleene–Rosser template $\Phi_{\mathrm{KR}}(p) = \neg\,\mathrm{Prov}_\tau(p)$, where $\mathrm{Prov}_\tau\colon \mathrm{Prop}_\tau \to \mathrm{Prop}_\tau$ is the canonical provability-predicate tail transformer of $\mathbf{Cat}_\tau$, the $\omega$-germ stabilised value depends on the behaviour of the internal provability iteration:*

(a) If $\mathrm{Prov}_\tau$ is tail-coherent at depth $k_0$ and the iteration oscillates with period 2 *(the Gödelian "diagonal Liar"), then* $[\![\phi]\!] = \mathrm{Both}$.

(b) If $\mathrm{Prov}_\tau$ fails to stabilise within the witness budget *(the classical "self-referential sentence is neither provable nor disprovable"), then* $[\![\phi]\!] = \mathrm{Neither}$.

*Case (b) is the $\tau$-topos's constructive analogue of Gödel's first incompleteness theorem: the self-referential sentence "this sentence is not provable" receives the truth value* Neither*, reflecting the* ontic *fact that no finite- witness provability certificate exists for it.*

*Proof.* Set $\Psi_{\mathrm{KR}} := \neg \circ \mathrm{Prov}_\tau$. In case (a), $\Psi_{\mathrm{KR}}$ acts on $B_\sigma(\mathbb{D})$ as a period-2 dynamical system (the negation forces the swap, and $\mathrm{Prov}_\tau$ is tail-coherent enough to preserve the single-lobe support). Theorem 7.7(c) then gives $[\![\phi]\!] = \mathrm{Both}$.

In case (b), $\mathrm{Prov}_\tau$ is not tail-coherent on the self-referential argument within the witness budget: the provability predicate applied to its own negated diagonal does not produce an NF-representable certificate in finitely many steps. This is the familiar Gödelian non-decidability, internal to $\mathbf{Cat}_\tau$. Theorem 7.7(d) then gives $[\![\phi]\!] = \mathrm{Neither}$.

The two cases correspond to the classical distinction between (i) $\phi$ as a *semantic* Liar, and (ii) $\phi$ as a *Gödelian independent* sentence. Both are expressible in $\mathbf{Cat}_\tau$; the first has a definite truth value at Both, the second at Neither. The $\tau$-topos therefore records the two incompleteness phenomena as two distinct truth sectors of the four-valued internal logic. □

**Remark 7.17** (Gödel's incompleteness as ontic non-stabilisation **[$\tau$-Effective]**)**.** Gödel's first incompleteness theorem [27, classically] asserts that arithmetic contains sentences neither provable nor disprovable from its axioms. In $\mathbf{Cat}_\tau$, such a sentence receives the truth value Neither $= 0 \in B_\sigma(\mathbb{D})$ — the ontic statement "no tail-coherence certificate within the witness budget." This is *not* an epistemic limitation ("we don't know whether $\phi$ is provable") but an ontic fact about the tail-transformer structure of $\mathrm{Prov}_\tau$: no NF-coded witness exists within the finite-witness discipline of $\mathbf{Cat}_\tau$. Gödel's theorem, interpreted in the $\tau$-topos, is the statement that $\mathrm{Prov}_\tau$ is not a total function on $\mathrm{SelfRef}_\tau$: some diagonal arguments land on Neither rather than on True or False, and this is a structural fact about the provability transformer, not a gap in the theory.

### 7.5 Ontic, not epistemic: the structural meaning of Both and Neither

A central claim of this paper is that the four truth values $\{\mathrm{Neither}, \mathrm{True}, \mathrm{False}, \mathrm{Both}\}$ of $\mathsf{Truth}_4 = B_\sigma(\mathbb{D})$ are *ontic* — genuine elements of the Boolean sublattice of the split-complex boundary algebra — rather than *epistemic* hedges on an underlying two-valued truth.

**Remark 7.18** (Ontic vs. epistemic: the distinction **[$\tau$-Effective]**)**.** An *epistemic* resolution treats Neither as "we don't know whether $p$ or $\neg p$" and Both as "we have incomplete information" — truth values function as second-order statements about a first-order binary truth (e.g. Kleene's three-valued logic [27]). An *ontic* resolution treats Neither and Both as first-order truth values: genuine algebraic elements of $B_\sigma(\mathbb{D}) \subset \mathbb{D}$, not placeholders.

(a) Both *is plenitude, not incompleteness.* Both $= 1 \in \mathbb{D}$ is the *multiplicative unit* of the split-complex algebra — the algebraic identity of $B_\sigma(\mathbb{D})$, the join $e_+ \vee e_-$, and the $\sigma$-fixed idempotent unit of the real axis $\mathbb{D}^\sigma = \mathcal{R}'_\partial \cdot 1$ (Remark 2.4). Stabilisation to Both is simultaneous support on both lobes — the Hegelian "unity of opposites" realised as a concrete algebraic identity.

(b) Neither *is non-support, not ignorance.* Neither $= 0 \in \mathbb{D}$ is the *zero element* — the additive identity annihilated by every multiplication. Stabilisation to Neither means no lobe-support: no $e_+$-sector witness, no $e_-$-sector witness, no Cauchy certificate within the finite-witness budget. This is an ontic algebraic fact about the tail transformer, not an epistemic confession.

The ontic interpretation is forced by the fact that $B_\sigma(\mathbb{D})$ is a *primary* algebraic object (the canonical Boolean sublattice of $\mathbb{D}$, existing prior to any question of truth) rather than a Boolean object derived from a pre-existing binary truth assignment. This is the deepest structural difference between the $\tau$-topos's resolution and the classical epistemic approaches.

**Remark 7.19** (Epistemic glosses as derived approximations **[$\tau$-Effective]**)**.** Informal glosses are *derived* from the ontic content, not definitional: True $= e_+$ (plus-lobe idempotent) glosses as "$p$ holds, $\neg p$ does not"; False $= e_-$ as the opposite; Both $= 1$ (joint idempotent unit) as "$p$ and $\neg p$ jointly hold"; Neither $= 0$ (zero, no support) as "no tail-coherence certificate." These are mnemonics; the *definitions* live in the algebra of $\mathbb{D}$.

**Remark 7.20** (Why four, and only four **[$\tau$-Effective]**)**.** Four-valuedness is algebraically forced: $B_\sigma(\mathbb{D})$ has exactly four elements by Hinge 4's uniqueness theorem [16, Thm. 1.8]. Three-valued logics cannot distinguish the two single-lobe atoms $e_+$, $e_-$ from an "intermediate" value and lack Both $= 1$ as the Liar's fixed point; five-or-more-valued logics introduce algebraically redundant sectors with no $B_\sigma(\mathbb{D})$-support. The $\tau$-topos is not free to choose a different truth cardinality.

### 7.6 Connection to Hinge 7's NF confluence

The classification theorem 7.7 is stated at **[$\tau$-Effective]** *modulo Hinge 7 canonical-address NF confluence* (Lemma 2.13). We now tabulate where exactly the Hinge-7 dependency enters.

**Remark 7.21** (Where NF confluence is used **[$\tau$-Effective]**)**.** Three steps in the construction of this section require the strong form of Lemma 2.13:

(a) *Well-definedness of* $[\![\Phi^n(s_\star)]\!]$. Each NF representative of $\Phi^n(s_\star)$ must be unique up to canonical-address equivalence, otherwise the sector projections $\pi_\pm$ would produce different images depending on the rewriting path. Weak confluence at the Hinge-5 level gives uniqueness up to $\sim$, which is enough for the $\sim$-tail equivalence class; strict confluence (Hinge 7) gives uniqueness at the level of the NF address itself, which is what promotes the classification to a deterministic algorithm.

(b) *Convergence of the $\omega$-germ limit.* The identification of a period-2 oscillation (case (c) of Theorem 7.7) with the algebraic sum $e_+ + e_-$ requires that the Cesàro average of the oscillating $\sim$-tail classes is itself an NF-representable element of $B_\sigma(\mathbb{D})$. Under weak confluence this holds up to $\sim$; under strict confluence it holds as a canonical address.

(c) *Distinguishing cases (c) and (d).* A sequence that looks non-stabilised within a given witness budget $k_0$ may turn out to be periodic at a deeper budget $k_0' > k_0$. Strict NF confluence guarantees that the distinction between *eventually periodic* and *persistently non- stabilised* is canonical (not budget-dependent): there exists a universal bound, controlled by the

NF-address depth of $\Phi$, beyond which the asymptotic behaviour is fixed.

Pending Hinge 7, the classification is *constructive and decidable in the pass-finite sense*: given a concrete NF code for $\Phi$ and a concrete witness budget $k_0$, one can algorithmically determine which of cases (a)–(d) applies, up to the weak confluence obstruction. Strict confluence upgrades this to a canonical classification.

**Remark 7.22** (The NF confluence conjecture: scope-tier summary **[$\tau$-Effective]**)**.** Theorem 7.7 and the four example theorems (7.10, 7.12, 7.14, 7.16) are all **[$\tau$-Effective]** *modulo Hinge 7 canonical-address NF confluence*. Upon certification of Hinge 7 (*Address Resolution, Not Calculation*, forthcoming), the scope promotes to **[$\tau$-Effective]** in the strict sense, and the classification becomes a theorem of the ambient $\tau$-framework without external hypothesis. The classical-topos embedding of Book II [9] then promotes further to **[Established]** by comparison with the classical Belnap–Dunn four-valued lattice, subject to a structural interpretation theorem that identifies $B_\sigma(\mathbb{D})$ with the classical four-element Boolean algebra of paraconsistent logic.

### 7.7 Comparison with classical resolutions

We conclude the section with a structural comparison of the $\omega$-germ stabilisation resolution against the principal classical approaches to semantic circularity. The comparison is labelled **[Established]** throughout: the classical resolutions are historical/structural facts from the philosophical-logic literature, and no proofs are claimed or re-derived here.

**Remark 7.23** (Comparison with Tarski's hierarchy **[Established]**)**.** Tarski's resolution restricts self-reference syntactically: truth predicates for an object language $\mathcal{L}_0$ live in a meta- language $\mathcal{L}_1$, whose truth predicate lives in $\mathcal{L}_2$, and so on [27]. The Liar $L = \neg L$ is simply *not well-formed* in any $\mathcal{L}_n$. The $\tau$-topos approach is *non-hierarchical*: self-reference is permitted as an NF-coded tail transformer (Definition 7.1) and the Liar acquires a definite truth value Both $= 1 \in B_\sigma(\mathbb{D})$ by Theorem 7.10. Structurally, the $\tau$-topos resolution is a categorical collapse of the Tarskian tower: a single elementary topos with a single four-valued truth space $\Omega_\tau = B_\sigma(\mathbb{D})$, in place of the would-be $\mathcal{L}_\omega$ union of all meta-language levels.

**Remark 7.24** (Comparison with three-valued logics **[Established]**)**.** Kleene's three-valued logic adds a single "undefined" value $u$ to $\{T, F\}$, and Łukasiewicz's three-valued logic analogously adds $\frac{1}{2}$ [27]; both resolve the Liar by $[\![L]\!] = u$. The three-valued resolution collapses Neither and Both into a single "undefined," losing the distinction between "no support" (ontic non-stabilisation) and "joint support" (paraconsistent unity). The four-valued $B_\sigma(\mathbb{D})$ separates them: the Liar is Both, Gödel's independent sentence is Neither, the Truth-teller is ambiguous. A three-valued logic cannot make this distinction; a four-valued logic over $B_\sigma(\mathbb{D})$ does so by algebraic force, since the four-atom cardinality is fixed by Hinge 4's uniqueness theorem for $\mathbb{D}$ [16, Thm. 1.6].

**Remark 7.25** (Comparison with Belnap–Dunn, Priest, Kripke **[Established]**)**.** The Belnap–Dunn four-valued logic [2] axiomatises exactly the four truth values $\{T, F, B, N\}$ with paraconsistent semantics; Priest's LP [27] is a closely related three-valued paraconsistent system; Kripke's fixed-point truth theory constructs a partial truth predicate by iterating a monotone jump operator on a three-valued lattice. $\mathsf{Truth}_4$ *coincides* with the Belnap–Dunn logic as a semantic object (same four values, same meet/join structure, same paraconsistent negation). What distinguishes the $\tau$-topos approach is the *derivation*: Belnap–Dunn axiomatises, Priest philosophically motivates, Kripke iterates over a transfinite set-theoretic universe. In $\mathbf{Cat}_\tau$ the four truth values are *derived* as the four atoms of the canonical Boolean sublattice $B_\sigma(\mathbb{D})$ of the split-complex boundary algebra $\mathbb{D}$ (Remark 2.3), via a uniqueness theorem about $\tau$-kernel structural constraints, and the Cauchy iteration is a countable, pass-finite procedure within the $\tau$-kernel finite-witness discipline — not a transfinite existence proof. The paraconsistent Belnap–Dunn logic is therefore *earned*, not stipulated.

**Remark 7.26** (Structural summary of the comparison **[Established]**)**.** Three axes distinguish the approaches: *(a) Self-reference permitted?* Tarski: no. Kleene, Łukasiewicz, Belnap–Dunn, Priest, Kripke, $\tau$-topos: yes. *(b) Truth-value cardinality.* Classical: 2; Kleene/Łukasiewicz/Priest LP/Kripke: 3; Belnap–Dunn/$\tau$-topos: 4. *(c) Axiomatic or derived?* All classical: axiomatic. $\tau$-topos: *derived* from the boundary algebra $\mathbb{D}$. The $\tau$-topos is, to our knowledge, the unique approach that is both four-valued paraconsistent *and* derives the four truth values from a primary algebraic structure — the central philosophical claim of this paper, anchored by Theorems 1.1, 1.2, 1.3, and 1.4.

### 7.8 Section summary

Theorem 7.7 is the paper's philosophical climax: it shows that every self-referential proposition in the $\tau$-topos admits a definite truth value in $B_\sigma(\mathbb{D}) = \mathsf{Truth}_4$, computed by the $\omega$-germ stabilisation of the Cauchy iteration $\Phi^n(s_\star)$. The four cases — True, False, Both, Neither — exhaust the asymptotic behaviours of the iteration and correspond to the four atoms of the Boolean sublattice. The Liar stabilises at Both $= 1$ (the algebraic idempotent unit); Curry admits two resolutions (paraconsistent-fixed-point Both or ontic-non-stabilisation Neither) depending on its NF admissibility; the Truth-teller is structurally ambiguous across all four atoms; Kleene–Rosser lands on Both or Neither in parallel with the Liar and Gödel respectively.

The resolution is *ontic, not epistemic*: the four truth values are the four algebraic elements of $B_\sigma(\mathbb{D}) \subset \mathbb{D}$, not hedges on an underlying binary truth. Both $= 1$ is the multiplicative unit of the split-complex algebra — the Hegelian unity of opposites realised algebraically; Neither $= 0$ is the additive identity — the ontic "no support" statement, not an epistemic confession of ignorance.

The classification is stated at **[$\tau$-Effective]** modulo Hinge 7's canonical-address NF confluence, and becomes a theorem of the ambient $\tau$-framework upon certification of Hinge 7. The comparison with classical resolutions (Tarski, Kleene, Belnap–Dunn, Priest, Kripke) shows that the $\omega$-germ approach is the unique resolution that is both four-valued paraconsistent *and* derives the four truth values from a primary algebraic structure. In a phrase: the four truth values of the $\tau$-topos are earned, not assumed.

Having established the internal-logic layer, we turn in §8 to the structural comparison of $\mathbf{Cat}_\tau$ with classical topos theory and the Grothendieck–topos embedding programme of Book II [9].

## 8. COMPARISON WITH CLASSICAL TOPOS THEORY

### 8.1 Elementary topoi in the Lawvere–Tierney sense

We begin with the classical notion of an elementary topos [26, 23], stated in the Lawvere–Tierney axiomatic form.

**Definition 8.1** (Classical elementary topos **[Established]**). *An* elementary topos *is a category $\mathcal{E}$ such that*

*(i) $\mathcal{E}$ has all finite limits;*

*(ii) $\mathcal{E}$ is cartesian closed (exponentials $Y^X$ exist for every pair of objects $X, Y$);*

*(iii) there is a* subobject classifier $\Omega \in \mathcal{E}$ *together with a mono $\top\colon 1 \rightarrowtail \Omega$ such that every monomorphism $m\colon S \rightarrowtail X$ arises as the pullback of $\top$ along a unique characteristic morphism $\chi_m : X \to \Omega$.*

*The* internal logic *of $\mathcal{E}$ is the first-order typed logic whose propositional fragment is the Heyting algebra* $\mathrm{Sub}(1) = \mathrm{Hom}(1, \Omega)$.

By Theorem 1.1, together with the constructions of §3, §4, and §5, the $\tau$-topos $\mathbf{Cat}_\tau$ satisfies Definition 8.1 clauses (i)–(iii). Explicitly:

- Finite limits are earned from $\mathsf{HolEnd}_\tau$ via the pre-Yoneda collapse of Hinge 5 [17].
- Exponentials $Y^X$ are realised as boundary-addressed representing objects for the hom-sheaf $\mathsf{HolEnd}_\tau(- \times X, Y)$ (see §3).
- The subobject classifier is $\Omega_\tau \cong B_\sigma(\mathbb{D})$ with characteristic morphism $\chi_\tau\colon\ \mathrm{Sub}_\tau \to \Omega_\tau$ (see §4).

*Hence* $\mathbf{Cat}_\tau$ *is an elementary topos in the Lawvere–Tierney sense.* The structural divergence from the classical picture enters only at the internal-logic layer, where the Heyting-algebra requirement of Definition 8.1 is replaced by a bilattice structure.

**Definition 8.2** (Elementary topos with paraconsistent internal logic **[$\tau$-Effective]**). *An* elementary topos with paraconsistent internal logic *is an elementary topos $\mathcal{E}$ in the sense of Definition 8.1 whose subobject classifier $\Omega$ is equipped with a bilattice structure $(\Omega, \leq_t, \leq_k, \wedge, \vee, \otimes, \oplus, \neg, \bot, \top)$ in place of the classical Heyting-algebra structure, and whose internal propositional calculus is interpreted over this bilattice rather than over a Heyting algebra.*

Definition 8.2 is a strict generalisation of Definition 8.1: every classical topos is a paraconsistent topos with trivial bilattice structure (collapsing $\leq_t$ and $\leq_k$ onto the classical order), while the converse fails precisely at the four-valued or larger bilattice cases.

**Proposition 8.3** ($\mathbf{Cat}_\tau$ is a paracontinuous elementary topos **[$\tau$-Effective]**). *The $\tau$-topos $\mathbf{Cat}_\tau$ is an elementary topos with paraconsistent internal logic in the sense of Definition 8.2, with $\Omega = \Omega_\tau = B_\sigma(\mathbb{D})$ and the bilattice structure $(B_\sigma(\mathbb{D}), \leq_t, \leq_k)$ of §5.*

*Proof sketch.* Theorem 1.1 supplies the elementary-topos structure of clauses (i)–(iii). Theorem 1.2 supplies the bilattice interpretation and its paraconsistent soundness. Together they witness Definition 8.2. □

### 8.2 Grothendieck topoi and the comparison functor

We next compare $\mathbf{Cat}_\tau$ with the classical Grothendieck construction of sheaves on a site [1, 22].

**Definition 8.4** (Grothendieck topos **[Established]**). *A* Grothendieck topos *is a category equivalent to* $\mathbf{Sh}(\mathcal{C}, J)$*, the category of sheaves on a small site* $(\mathcal{C}, J)$ *with respect to a Grothendieck topology* $J$*. Equivalently (Giraud's theorem), it is a locally presentable, cocomplete cartesian-closed category whose finite limits commute with filtered colimits and which admits a generating small full subcategory.*

Classical Grothendieck topoi inherit their internal logic from the Heyting-algebra structure of $\Omega = \mathrm{Sub}(1)$; this is a theorem (not an axiom) of the Giraud construction [26, Ch. III]. This is precisely where $\mathbf{Cat}_\tau$ diverges from the classical Grothendieck setting.

**Remark 8.5** (Why $\mathbf{Cat}_\tau$ is not a classical Grothendieck topos **[$\tau$-Effective]**). The $\tau$-topos $\mathbf{Cat}_\tau$ is not a Grothendieck topos in the sense of Definition 8.4 for three structural reasons:

1. The subobject classifier $\Omega_\tau = B_\sigma(\mathbb{D})$ is four-valued rather than two-valued; a sheaf topos on any site inherits $\Omega$ from the base, and a two-valued base cannot furnish $B_\sigma(\mathbb{D})$ directly.
2. The internal logic is paraconsistent (bilattice-valued), not Heyting. Every classical Grothendieck topos is Heyting by [26, Prop. V.1].
3. The pre-Yoneda collapse of Hinge 5 is specific to $\omega$-germ stabilisation on the countable profinite boundary ring $\mathcal{R}'_\partial$; it is not the Grothendieck-sheafification $\mathbf{PSh}_\tau \to \mathbf{Sh}_\tau$ of a classical site.

Nevertheless, a canonical bridge exists, constructed next.

**Definition 8.6** (Boundary-addressable Grothendieck site $\mathcal{S}$ **[$\tau$-Effective]**). *Let $\mathcal{S}$ be the small category whose objects are boundary-addressable carriers $(X, \alpha_X)$ where $\alpha_X$ is a countable address scheme on $X$, and whose morphisms are address-preserving $\omega$-germ morphisms modulo tail-equivalence. Equip $\mathcal{S}$ with the* pre-Yoneda-collapse coverage*: a family $\{f_i : U_i \to X\}_{i \in I}$ covers $X$ iff the induced morphism $\coprod_i U_i \to X$ is a tail-equivalence after $\omega$-germ stabilisation. Write $\mathbf{Sh}_\tau(\mathcal{S})$ for the resulting Grothendieck topos of sheaves on $(\mathcal{S}, J_\tau)$.*

**Theorem 8.7** (Comparison functor $\Psi$ **[$\tau$-Effective]**). *There is a canonical comparison functor*

$$\Psi : \mathbf{Cat}_\tau \longrightarrow \mathbf{Sh}_\tau(\mathcal{S}) \tag{38}$$

*which preserves finite limits and exponentials, and which sends the subobject classifier $\Omega_\tau$ to the classical (two-valued) subobject classifier of $\mathbf{Sh}_\tau(\mathcal{S})$ by the quotient map $\pi_{\mathrm{cl}} : B_\sigma(\mathbb{D}) \twoheadrightarrow \{0, 1\}$ defined by $\pi_{\mathrm{cl}}(\mathrm{Neither}) = 0$, $\pi_{\mathrm{cl}}(\mathrm{True}) = 1$, $\pi_{\mathrm{cl}}(\mathrm{False}) = 0$, $\pi_{\mathrm{cl}}(\mathrm{Both}) = 1$.*

*Proof sketch.* On objects, $\Psi$ sends a boundary-addressable carrier $X$ in $\mathbf{Cat}_\tau$ to the representable sheaf $y(X) = \mathrm{Hom}_{\mathcal{S}}(-, X)$ in $\mathbf{Sh}_\tau(\mathcal{S})$. On morphisms, $\Psi$ sends an $\omega$-germ morphism to its post-stabilisation representative in the representable sheaf. Finite-limit preservation is immediate from the Yoneda lemma; exponential preservation follows from the pre-Yoneda collapse of Hinge 5 and the compatibility of $\omega$-germ stabilisation with hom-internalisation. The subobject classifier is *not* preserved as a whole (the four-valued $\Omega_\tau$ maps onto the two-valued $\Omega_{\mathbf{Sh}_\tau}$), but the collapse rule $\pi_{\mathrm{cl}}$ is functorial and respects $\wedge, \vee, \top, \bot$ on the Boolean sublattice image. □

The functor $\Psi$ of Theorem 8.7 makes precise the sense in which the classical Grothendieck topos view "sees" only the Boolean shadow $\{0, 1\}$ of the four-valued structure, collapsing the paraconsistent distinction Neither $\not\leftrightarrow \bot$ and Both $\not\leftrightarrow \top$ of $B_\sigma(\mathbb{D})$.

### 8.3 Heyting algebras versus bilattices

We now articulate the key algebraic distinction between classical and $\tau$-topos internal logic.

**Definition 8.8** (Heyting algebra **[Established]**). *A* Heyting algebra *is a distributive bounded lattice* $(H, \leq, \wedge, \vee, \top, \bot)$ *equipped with a binary operation* $\to$ *(*relative pseudocomplement*) satisfying the adjunction*

$$p \wedge q \leq r \quad \Longleftrightarrow \quad q \leq (p \to r) \qquad \text{for all } p, q, r \in H. \tag{39}$$

*Equivalently, H is the propositional fragment of intuitionistic logic.*

The Boolean lattice $\{0, 1\}$ of classical propositional logic is the initial Heyting algebra satisfying the law of excluded middle $p \vee \neg p = \top$. Classical (Boolean) topoi have $\Omega = \{0, 1\}$; general classical Grothendieck topoi have $\Omega$ a more general Heyting algebra (a frame).

**Remark 8.9** (Why $B_\sigma(\mathbb{D})$ is not a Heyting algebra **[$\tau$-Effective]**). The four-element bilattice $B_\sigma(\mathbb{D}) = \{\text{Neither}, \text{True}, \text{False}, \text{Both}\}$ of $\mathsf{Truth}_4$ is *not* a Heyting algebra. Two explicit failures:

- *Non-contradiction fails*: $p \wedge \neg p = \text{Both} \neq \bot = \text{Neither}$ when $[\![p]\!] = \text{Both}$. In a Heyting algebra, $p \wedge \neg p \leq \bot$ always holds.
- *The Heyting adjunction* (39) *fails*: there is no binary operation $\to_{\mathrm{H}}$ on $B_\sigma(\mathbb{D})$ satisfying (39) for all $p, q, r$. Witness: take $p = \text{Both}$, $q = \text{True}$, $r = \text{False}$; any attempted $p \wedge q = \text{True} \leq_t \text{False} = r$ fails in the truth order, but the putative pseudocomplement $p \to_{\mathrm{H}} r = \text{Both} \to_{\mathrm{H}} \text{False}$ must simultaneously entail True $\leq_t$ something matching False, which no element of $B_\sigma(\mathbb{D})$ provides.

Consequently $B_\sigma(\mathbb{D})$ does *not* support an internal intuitionistic logic. What it *does* support is the stronger bilattice structure described next.

**Definition 8.10** (Bilattice **[Established]**). *A* bilattice *is a structure* $(B, \leq_t, \leq_k)$ *where* $(B, \leq_t)$ *and* $(B, \leq_k)$ *are both bounded lattices, and the meet/join operations of each order distribute appropriately with those of the other. The* truth order $\leq_t$ *encodes* degree of truth *(*$\text{False} \leq_t \text{Neither}, \text{Both} \leq_t \text{True}$*); the* knowledge order $\leq_k$ *encodes* degree of information *(*$\text{Neither} \leq_k \text{True}, \text{False} \leq_k \text{Both}$*). See [2, 27] for the foundational development.*

**Proposition 8.11** (Bilattices properly generalise Heyting algebras **[Established]**). *Every Heyting algebra H extends to a bilattice* $H \times H^{\mathrm{op}}$ *(the classical "squared" bilattice construction), and conversely every bilattice with a single "collapsed"* $\leq_t = \leq_k$ *is a distributive lattice. Bilattices are a strictly larger class:* $B_\sigma(\mathbb{D})$ *is a bilattice but* not *a Heyting algebra (Remark 8.9).*

**Remark 8.12** (Structurally grounded bilattice **[$\tau$-Effective]**). Unlike most bilattice-valued logics in the literature [2, 27], where the four (or more) truth values are *axiomatically* stipulated, the bilattice structure of $B_\sigma(\mathbb{D})$ in $\mathbf{Cat}_\tau$ is *earned* from the split-complex boundary algebra $\mathbb{D}$. Specifically:

- $\{0, e_+, e_-, 1\}$ is the $\sigma$-equivariant idempotent sublattice of the unique (up to canonical isomorphism) split-complex algebra $\mathbb{D} = \mathcal{R}'_\partial[j]/(j^2 - 1)$ (Hinge 4, [16]).
- The truth order $\leq_t$ arises from the multiplicative ordering of idempotents.
- The knowledge order $\leq_k$ arises from the $\omega$-germ stabilisation partial order (Hinge 5, [17]).

Thus the four truth values, their orders, and their interactions are *theorems* about $\mathbb{D}$, not axiomatic choices. This structural grounding is what distinguishes the $\tau$-topos from earlier paraconsistent topos proposals.

### 8.4 The classical-topos subquotient

Having identified the structural divergence, we now construct the classical shadow of $\mathbf{Cat}_\tau$ and show it recovers a known Grothendieck topos.

**Definition 8.13** (Classical-topos equivalence $\sim_{\mathrm{cl}}$ **[$\tau$-Effective]**). *Define the* classical-topos equivalence $\sim_{\mathrm{cl}}$ *on* $B_\sigma(\mathbb{D})$ *by*

$$\text{Neither} \sim_{\mathrm{cl}} \text{False}, \qquad \text{True} \sim_{\mathrm{cl}} \text{Both}\,. \tag{40}$$

*The quotient $B_\sigma(\mathbb{D})/{\sim_{\mathrm{cl}}} \cong \{0,1\}$ is the classical two-element Boolean lattice under the map $\pi_{\mathrm{cl}}$ of Theorem 8.7. Lift $\sim_{\mathrm{cl}}$ to $\mathbf{Cat}_\tau$ pointwise: two morphisms $f, g\colon X \to Y$ are classically equivalent iff $\chi_\tau(f) \sim_{\mathrm{cl}} \chi_\tau(g)$ for all subobject queries.*

**Theorem 8.14** (Classical-topos subquotient **[$\tau$-Effective]**). *The quotient category $\mathbf{Cat}_\tau^{\mathrm{cl}} := \mathbf{Cat}_\tau/{\sim_{\mathrm{cl}}}$ is a classical elementary topos in the sense of Definition 8.1, with two-valued subobject classifier $\Omega^{\mathrm{cl}} = \{0,1\}$ and Heyting internal logic. Moreover, $\mathbf{Cat}_\tau^{\mathrm{cl}}$ is equivalent to the Grothendieck topos $\mathbf{Sh}_\tau(\mathcal{S}_{\mathrm{triv}})$ of boundary-addressable sheaves on the trivial site $\mathcal{S}_{\mathrm{triv}}$ whose only covering families are the identity covers.*

*Proof sketch.* That $\mathbf{Cat}_\tau^{\mathrm{cl}}$ is cartesian closed with finite limits follows from Theorem 1.1 together with the standard fact that a topos quotient by a congruence respecting finite limits and exponentials is again a topos (see [23, Ch. A4]). The classical subobject classifier is $\pi_{\mathrm{cl}}(\Omega_\tau) = \{0,1\}$, which is the initial Boolean algebra and thus the classifier of $\mathbf{Sh}(\mathcal{S}_{\mathrm{triv}})$. The equivalence with $\mathbf{Sh}_\tau(\mathcal{S}_{\mathrm{triv}})$ follows from the Yoneda embedding: the trivial-site sheaf topos is equivalent to the category of presheaves on $\mathcal{S}_{\mathrm{triv}}$, and $\mathbf{Cat}_\tau^{\mathrm{cl}}$ presents the same universal property via the pre-Yoneda collapse modulo $\sim_{\mathrm{cl}}$. □

**Corollary 8.15** (Classical facts have $\tau$-analogues **[$\tau$-Effective]**). *Every classical-topos theorem (Lawvere–Tierney factorisation, Diaconescu's theorem, the Freyd cover, etc.) has a direct analogue in $\mathbf{Cat}_\tau$ via passage through $\mathbf{Cat}_\tau^{\mathrm{cl}}$, where the analogue is obtained by adjoining the extra bilattice structure on the four-valued subobject classifier. The classical statement is recovered as the $\sim_{\mathrm{cl}}$-image; the $\tau$-statement refines it by tracking the* Neither/Both *distinction.*

Corollary 8.15 is the content-level justification for calling $\mathbf{Cat}_\tau$ a "genuine generalisation": the classical content is *recoverable in full*, and the $\tau$-extension adds further structure (the Neither and Both sectors) that was not available in the classical world.

### 8.5 Homotopy type theory connection

We briefly sketch the connection to homotopy type theory [31] and univalent foundations, in the spirit of Hinge 5's ontological-primary framing.

**Remark 8.16** ($\tau$-topos as paraconsistent univalent host **[Conjectural]**). The type-theoretic semantics of $\mathbf{Cat}_\tau$ is naturally homotopy-theoretic:

- Boundary-addressable carriers $X \in \mathbf{Cat}_\tau$ correspond to *types* in the sense of HoTT.
- Propositional equality $x \equiv y$ on a carrier corresponds to $\omega$-germ tail-equivalence $x \sim y$.
- The univalence axiom $(X \simeq Y) \simeq (X = Y)$ would correspond to the statement that $\omega$-germ stabilisation produces a *canonical* path between any two judgmentally equal terms (the pre-Yoneda collapse furnishes such a canonical path at the level of representable objects).
- The four-valued nature of $\Omega_\tau$ extends HoTT's propositional hierarchy by a *paraconsistent flavour*: the Both sector admits self-referential propositions that stabilise to the idempotent unit; the Neither sector admits propositions that have not yet stabilised.

A full HoTT treatment — including univalence, higher inductive types, and $\infty$-categorical refinements — is deferred to Book II [9]. We merely observe here that $\mathbf{Cat}_\tau$ is a natural host for univalent foundations with paraconsistent flavour, and that standard classical HoTT is recovered as the $\sim_{\mathrm{cl}}$-quotient.

The conjectural status of Remark 8.16 is essential: we do not here claim that $\mathbf{Cat}_\tau$ is a model of HoTT (nor that such a model exists within the countable addressable setting); we claim only that the type-theoretic shape of $\mathbf{Cat}_\tau$ matches HoTT's, with the bilattice structure as a proper enrichment.

### 8.6 Categorical logic foundations

We close the substantive comparisons with three brief positionings relative to foundational work in categorical logic and constructivism.

**Johnstone's *Elephant*..** [23] hints, in Part A Chapter 4, at a programme of elementary topoi with *generalised truth-value objects* — topoi whose subobject classifier is a non-Boolean bounded lattice with further algebraic structure. The *Elephant* does not fully develop a bilattice-valued theory, but Johnstone's framework of classifying topoi for geometric theories already anticipates the possibility of multi-valued truth in elementary settings. The $\tau$-topos $\mathbf{Cat}_\tau$ instantiates this programme with a structurally earned four-valued bilattice classifier, and (by Theorem 8.14) reduces to the classical case via a concrete reflective quotient.

**Bishop–Bridges constructive analysis..** [3] establishes that mathematical analysis admits a fully constructive development without the law of excluded middle, using explicit constructions and decidable procedures. Three features of $\mathbf{Cat}_\tau$ align directly with the Bishop–Bridges programme:

- *Countability and addressability*: every object, morphism, and proposition in $\mathbf{Cat}_\tau$ is countable and boundary-addressable via the profinite ring $\mathcal{R}'_\partial$.
- *Decidable procedures*: the classification $[\![p]\!] \in \{\text{Neither, True, False, Both}\}$ of Theorem 1.3 is a finite-witness decidable procedure, in full agreement with Bishop's constructivism.
- *No LEM*: the law of excluded middle fails at the Neither sector, and the $\tau$-topos does not presuppose it.

$\mathbf{Cat}_\tau$ is thus, in a precise sense, a Bishop-constructive elementary topos with paraconsistent enrichment.

**Troelstra–van Dalen constructivism..** [30] develops classical (Brouwerian) constructivism from an intuitionistic-logic base. The paraconsistent extension we supply in $\mathbf{Cat}_\tau$ is *orthogonal* to — not in conflict with — classical intuitionistic constructivism: the $\sim_{\text{cl}}$-quotient of $\mathbf{Cat}_\tau$ onto $\mathbf{Cat}_\tau^{\text{cl}}$ recovers the classical intuitionistic setting, and the bilattice extension adds information-order structure without removing any intuitionistic content. In particular, the double-negation translation of classical logic into intuitionistic logic (Gödel–Gentzen–Kolmogorov) extends to a bilattice-double-negation translation internalised in $\mathbf{Cat}_\tau$; we do not develop this here.

### 8.7 Summary comparison table

Table 1 summarises the structural comparison between classical topos theory and the $\tau$-topos $\mathbf{Cat}_\tau$. Features common to both sides are attributes of any elementary topos; features distinguishing the sides are the sites of proper generalisation.

**Table 1.** Comparison of classical topos theory and the $\tau$-topos $\mathbf{Cat}_\tau$. The left column gives the feature; the middle column its classical realisation; the right column its $\tau$-topos realisation. The classical side is standard [26, 23]; the $\tau$-topos side is earned throughout §§2–7 of this paper.

| Feature | Classical topos | $\tau$-topos $\mathbf{Cat}_\tau$ |
|---|---|---|
| Finite limits | Axiomatic | Earned from $\mathsf{Hol}_\tau$ |
| Exponentials | Axiomatic | Earned via pre-Yoneda collapse |
| Subobject classifier | $\{0, 1\}$ (two-valued) | $B_\sigma(\mathbb{D}) = \{\text{Neither, True, False, Both}\}$ |
| Internal logic | Heyting algebra | Paraconsistent bilattice |
| Law of excluded middle | Holds (in Boolean topoi) | Fails at Neither |
| Law of non-contradiction | Holds | Fails at Both |
| Explosion ($p, \neg p \vdash q$) | Classical: yes | Fails (paraconsistent) |
| Cardinality | Typically uncountable | Countable (profinite addressable) |
| Source of truth values | Axiomatic stipulation | Earned from $\mathbb{D}$ algebra |
| Internal language | First-order typed | First-order typed (four-valued) |
| Classical content | Native | Recoverable via $\sim_{\text{cl}}$-quotient |
| Self-reference | Tarski hierarchy, restricted | $\omega$-germ stabilised in $B_\sigma(\mathbb{D})$ |

The pattern visible in Table 1 is uniform: in each row where classical and $\tau$-topos sides differ, the $\tau$-topos side is a *proper enrichment* that reduces to the classical case under the quotient $\sim_{\text{cl}}$ of Definition 8.13.

### 8.8 Reflective embedding of the classical topos

We close the section with the structural theorem promised in the introduction: the classical topos $\mathbf{Cat}_\tau^{\text{cl}}$ sits inside $\mathbf{Cat}_\tau$ as a reflective subcategory, and the $\tau$-extension is a genuine (that is, non-trivial) generalisation.

**Theorem 8.17** (Classical embedding **[$\tau$-Effective]**). *There is a full and faithful embedding*

$$\Psi: \mathbf{Cat}_\tau^{\mathrm{cl}} \hookrightarrow \mathbf{Cat}_\tau \tag{41}$$

*realising the classical topos* $\mathbf{Cat}_\tau^{\mathrm{cl}}$ *of Theorem 8.14 as a* reflective subcategory *of the* $\tau$*-topos* $\mathbf{Cat}_\tau$*. Concretely:*

*(a) The left adjoint (classical-quotient)* $Q: \mathbf{Cat}_\tau \to \mathbf{Cat}_\tau^{\mathrm{cl}}$ *is the quotient functor collapsing* $B_\sigma(\mathbb{D}) \to \{0, 1\}$ *via* $\pi_{\mathrm{cl}}$ *pointwise.*

*(b) The right adjoint (classical-embedding)* $\Psi$ *sends each classical-topos object to its "sharp" four-valued lift, defined so that only* True *and* False *sectors are used and* Neither/Both *are reserved for non-classical objects.*

*(c) The composition* $Q \circ \Psi = \mathrm{id}_{\mathbf{Cat}_\tau^{\mathrm{cl}}}$ *(the classical topos is recovered on the nose).*

*(d) The unit* $\eta_X: X \to \Psi(Q(X))$ *is the* $\sim_{\mathrm{cl}}$*-collapse natural transformation; it fails to be invertible precisely on the objects using the* Neither *or* Both *sectors.*

*Proof sketch. Fullness and faithfulness of* $\Psi$*.* Because $\Psi$ lifts a classical object to use only the $\{\text{True, False}\}$ sub-bilattice of $B_\sigma(\mathbb{D})$, the hom-sets are preserved isomorphically: any $\tau$-morphism between sharp lifts is already a classical-topos morphism.

*Adjunction.* The bijection $\mathrm{Hom}_{\mathbf{Cat}_\tau}(X, \Psi(Y)) \cong \mathrm{Hom}_{\mathbf{Cat}_\tau^{\mathrm{cl}}}(Q(X), Y)$ is witnessed by $\pi_{\mathrm{cl}}$ on the codomain side and by the sharp lift on the domain side. Naturality in $X$ and $Y$ follows from the functoriality of $\pi_{\mathrm{cl}}$ and the $\sim_{\mathrm{cl}}$-invariance of the classical-topos hom-sets.

*Reflectivity.* $Q \circ \Psi = \mathrm{id}$ because $\pi_{\mathrm{cl}}$ collapses $\{\text{True, False}\}$ onto $\{1, 0\}$ bijectively and the sharp lift uses only the $\{\text{True, False}\}$ sectors.

*Non-invertibility of* $\eta$ *on non-classical objects.* Any object $X$ using the Neither or Both sector has $\eta_X: X \to \Psi(Q(X))$ collapse the four-valued structure to two-valued, and the classical lift cannot recover Neither/Both; hence $\eta_X$ is not invertible. □

**Corollary 8.18** ($\mathbf{Cat}_\tau$ is a proper generalisation **[$\tau$-Effective]**). *The* $\tau$*-topos* $\mathbf{Cat}_\tau$ *is a proper generalisation of classical topos theory: the classical theory is recovered as a reflective subcategory (Theorem 8.17), and the paraconsistent enrichment corresponds to objects outside the reflective image — precisely those with non-trivial* Neither/Both *sector structure.*

**Remark 8.19** (Scope and caveats **[$\tau$-Effective]**). Three caveats are appropriate.

1. Theorem 8.17 is **[$\tau$-Effective]** because the full adjunction is conditional on Hinge 7's canonical-address NF confluence — specifically, the sharp lift must respect the unique canonical address on each classical object, and this is a Hinge 7 theorem.
2. The classical topos $\mathbf{Cat}_\tau^{\mathrm{cl}}$ is a classical Grothendieck topos (Theorem 8.14) whose underlying site is the trivial-covering boundary-addressable site $\mathcal{S}_{\mathrm{triv}}$; richer Grothendieck topologies (the étale, Zariski, or flat-coverage analogues) are deferred to Book II [9].
3. The HoTT connection of §8.5 remains **[Conjectural]**; any upgrade to **[$\tau$-Effective]** requires an explicit univalence-compatible model construction, also deferred to Book II.

In summary, the $\tau$-topos $\mathbf{Cat}_\tau$ occupies a position in the landscape of categorical logic that is:

- A *Lawvere–Tierney elementary topos* (by Proposition 8.3).
- *Not* a classical Grothendieck topos, but related to one via the comparison functor $\Psi$ of Theorem 8.7.
- A *proper extension* of its own classical subquotient $\mathbf{Cat}_\tau^{\mathrm{cl}}$, related reflectively via Theorem 8.17.
- A *Bishop-constructive* topos with a paraconsistent bilattice classifier, orthogonal to but compatible with Troelstra–van Dalen constructivism.
- A *natural host* for a paraconsistent variant of univalent foundations (deferred to Book II).

These five positions — all earned, not axiomatic — jointly justify regarding $\mathbf{Cat}_\tau$ as a genuine generalisation of classical topos theory, structurally grounded in the split-complex boundary algebra and the earned categorical machine of the Panta Rhei hinge bundle.

## 9. LEAN ROADMAP AND REGISTRY ENTRIES

### 9.1 Lean roadmap (detailed)

The planned Lean 4 formalisation lives in `TauLib.BookII.Topos` (per the $\iota_\tau$/PR-II alignment), with module structure:

- `TauTopos.lean` — the underlying category $\mathbf{Cat}_\tau$.
- `Truth4.lean` — the four-valued lattice $B_\sigma(\mathbb{D})$ with its paraconsistent operations.
- `SubobjectClassifier.lean` — $\Omega_\tau$ and $\chi_\tau$, with the universal property verified on the four truth-sectors.
- `ExponentialObjects.lean` — $Y^X$ via pre-Yoneda collapse. Depends on `TauLib.BookII.Holomorphy.HolEnd` and (forthcoming) `TauLib.BookI.Addressability` (Hinge 7 scope).
- `Circularity.lean` — constructive stabilisation of self-referential propositions.

### 9.2 Registry IDs

**Remark 9.1** (Registry IDs **[$\tau$-Effective]**). The five main theorems of this paper are registered in `registry/book2_registry.jsonl` as II.T66 (topos structure), II.T67 (paraconsistent soundness), II.T68 (circularity resolution), II.T69 (Both $= 1$), II.T70 (Hinge 6 integration), continuing the Hinge-5 range II.T57–T65. Registration follows peer-panel certification.

## 10. CONCLUSION AND FORWARD LINKS

The $\tau$-topos $\mathbf{Cat}_\tau$ presented here realises a genuinely *categorical* resolution of the classical paradoxes of self- reference and paraconsistent logic. Its four truth values $\{\text{Neither, True, False, Both}\} = \{0, e_+, e_-, 1\}$ are not axiomatic stipulations but earned theorems of the split-complex boundary algebra of Hinge 4 and the earned categorical machine of Hinge 5.

- **Hinge 7** (forthcoming) — canonical addressability via genealogical DAG and Cayley word metric; supplies the NF confluence theorem on which this paper's exponential-object construction ultimately depends.
- **Book II** [9] — classical-topos embedding, Grothendieck-topology comparison, sheaf-semantics extension.
- **Book III** [10] — applications to physical modal logic and the categorical semantics of measurement.
- **Books IV–VII** [11, 12, 13, 14] — applications in quantum mechanics, gravity, biology, and metaphysics where paraconsistent semantic circularity plays a structural role.

### Acknowledgments

The authors thank the Panta Rhei research community for feedback on earlier drafts, the Lean 4 and `mathlib` communities for the formal-logic infrastructure that underwrites this paper's planned mechanisation, and the coordinated development of Hinges 1–5 whose structural objects this paper now lifts into a genuine elementary topos.

### Data and code availability

The source repository for the paper bundle is at `https://panta-rhei.site/papers/tau-topos`. Planned Lean 4 artefacts for the main theorems will appear in `TauLib.BookII.Topos` (see §9).

## REFERENCES

[1] Michael Artin, Alexander Grothendieck, and Jean-Louis Verdier, *Théorie des topos et cohomologie étale des schémas (sga 4), tome 1*, Lecture Notes in Mathematics, vol. 269, Springer, Berlin, 1972.

[2] Nuel D. Belnap, *A useful four-valued logic*, Reidel, Dordrecht, 1977.

[3] Errett Bishop and Douglas Bridges, *Constructive analysis*, Grundlehren der mathematischen Wissenschaften, vol. 279, Springer, Berlin, 1985.

[4] Francesco Catoni, Dino Boccaletti, Roberto Cannata, Vincenzo Catoni, Enrico Nichelatti, and Paolo Zampetti, *The mathematics of minkowski space-time*, Frontiers in Mathematics, Birkhäuser, Basel, 2008.

[5] Thorsten Fuchs and Anna-Sophie Fuchs, *Address resolution, not calculation: The genealogical DAG, the Cayley metric, and why arithmetic in category $\tau$ has no equations*, 2026, Panta Rhei Research, `https://panta-rhei.site/papers/address-resolution`. Hinge 7 of the Panta Rhei bundle.

[6] ———, *The hyperfactorization theorem: Unique tower-atom decomposition in ZFC and category $\tau$*, 2026, Panta Rhei Research, https://panta-rhei.site/papers/hyperfactorization. Hinge 1 of the Panta Rhei bundle.

[7] ———, *The master constant $\iota_\tau$: A structural derivation of $\iota_\tau = 2/(\pi + e)$*, 2026, Panta Rhei Research, https://panta-rhei.site/papers/iota-tau. Hinge 3 of the Panta Rhei bundle.

[8] ———, *Panta rhei, Book I: Categorical Foundations*, 2nd ed., Panta Rhei Research, 2026, ISBN 979-8-3462-1987-6.

[9] ———, *Panta rhei, Book II: Categorical Holomorphy*, 2nd ed., Panta Rhei Research, 2026, ISBN 979-8-3462-1988-3.

[10] ———, *Panta rhei, Book III: Categorical Spectrum*, 2nd ed., Panta Rhei Research, 2026, ISBN 979-8-3462-1989-0.

[11] ———, *Panta rhei, Book IV: Categorical Microcosm*, 2nd ed., Panta Rhei Research, 2026, Volume IV of the 7-book Panta Rhei series; ISBN forthcoming.

[12] ———, *Panta rhei, Book V: Categorical Macrocosm*, 2nd ed., Panta Rhei Research, 2026, Volume V of the 7-book Panta Rhei series; ISBN forthcoming.

[13] ———, *Panta rhei, Book VI: Categorical Life*, 2nd ed., Panta Rhei Research, 2026, Volume VI of the 7-book Panta Rhei series; ISBN forthcoming.

[14] ———, *Panta rhei, Book VII: Categorical Metaphysics*, 2nd ed., Panta Rhei Research, 2026, Volume VII of the 7-book Panta Rhei series; ISBN forthcoming.

[15] ———, *The prime polarity theorem: A ZFC proof and the $\tau$-framework isomorphism*, 2026, Panta Rhei Research, https://panta-rhei.site/papers/prime-polarity. Hinge 2 of the Panta Rhei bundle.

[16] ———, *The split-complex boundary algebra $\mathbb{D}$: Canonical uniqueness, countable profinite construction, and the four-atom generator dictionary*, 2026, Panta Rhei Research, https://panta-rhei.site/papers/boundary-algebra. Hinge 4 of the Panta Rhei bundle.

[17] ———, *$\tau$-Holomorphy on the boundary algebra: $\omega$-Germ transformers, the wave-equation Cauchy–Riemann, and the earned categorical machine*, 2026, Panta Rhei Research, https://panta-rhei.site/papers/holomorphy-first. Hinge 5 of the Panta Rhei bundle.

[18] ———, *The $\tau$-kernel as foundational architecture: Ontic identity, linear discipline, and the star-autonomous path*, 2026, Panta Rhei Research, https://panta-rhei.site/papers/kernel-foundation. Hinge 8 (foundational anchor) of the Panta Rhei bundle; may be read as an entry point (Hinge 0) or as a capstone.

[19] ———, *TauLib: A Lean 4 library for the Panta Rhei research program*, 2026, Available at https://panta-rhei.site/verify/.

[20] Robert C. Gunning and Hugo Rossi, *Analytic functions of several complex variables*, Prentice-Hall, Englewood Cliffs, NJ, 1965.

[21] Lars Hörmander, *An introduction to complex analysis in several variables*, 3 ed., North-Holland, Amsterdam, 1990.

[22] Peter T. Johnstone, *Topos theory*, Academic Press, London, 1977.

[23] ———, *Sketches of an elephant: A topos theory compendium, vol. 1*, Oxford University Press, Oxford, 2002.

[24] Steven G. Krantz, *Function theory of several complex variables*, 2 ed., AMS Chelsea Publishing, American Mathematical Society, Providence, RI, 1992.

[25] Lean Community, *The Lean 4 theorem prover*, Software and documentation, 2024, https://leanprover.github.io/.

[26] Saunders Mac Lane and Ieke Moerdijk, *Sheaves in geometry and logic: A first introduction to topos theory*, Universitext, Springer, New York, 1992.

[27] Graham Priest, *An introduction to non-classical logic: From if to is*, 2 ed., Cambridge University Press, Cambridge, 2008.

[28] R. Michael Range, *Holomorphic functions and integral representations in several complex variables*, Graduate Texts in Mathematics, vol. 108, Springer, New York, 1986.

[29] Garret Sobczyk, *The hyperbolic number plane*, The College Mathematics Journal **26** (1995), no. 4, 268–280.

[30] A. S. Troelstra and Dirk van Dalen, *Constructivism in mathematics: An introduction, vol. i*, North-Holland, Amsterdam, 1988.

[31] Univalent Foundations Program, *Homotopy type theory: Univalent foundations of mathematics*, Institute for Advanced Study, Princeton, NJ, 2013, Available online at `https://homotopytypetheory.org/book/`.

[32] I. M. Yaglom, *A simple non-euclidean geometry and its physical basis*, Springer, New York, 1979.