HINGE 5 PREPRINT · V1.0 · APRIL 2026

τ-Holomorphy on the Boundary Algebra

ω-germ transformers, the wave-equation Cauchy–Riemann, and the earned categorical machine

Thorsten Fuchs • Anna-Sophie Fuchs

Correspondence: thorsten@panta-rhei.site

April 2026

DOI: 10.5281/zenodo.19818852

ABSTRACT

We establish τ-holomorphy as the *ontological primary* of Category τ, prior to any notion of "mapping," "function," "tuple," or "Cartesian product": a τ-holomorphic map is defined directly as a certified ω-germ transformer on the boundary algebra, and the entire categorical apparatus (composition, identity, associativity, functoriality) is then derived as earned theorem. Building on the Panta Rhei Boundary Algebra paper [?] (Hinge 4), which fixes the unique τ-admissible scalar algebra $\mathbb{D} = \mathcal{R}'_{\partial}[j]/(j^2 - 1)$, we prove: (i) a *legitimacy theorem* for why ω-germ transformers deserve the name "holomorphic"; (ii) an *earned scalar codomain* theorem forcing $\mathbb{D}$ as the unique codomain compatible with the boundary-transformer axioms; (iii) the *wave-equation Cauchy–Riemann* theorem, establishing that the components of a $\mathbb{D}$-holomorphic function satisfy the hyperbolic wave equation $\partial_t^2 f = \partial_x^2 f$ rather than the elliptic Laplace equation; (iv) a *diagonal-discipline* theorem, identifying the No-Cartesian-Product principle as the exact structural obstruction that protects $\mathbb{D}$ from the idempotent collapse that would force it toward elliptic $\mathbb{C}$; (v) an *earned categorical machine* theorem deriving composition, identity, associativity, and functoriality as theorems (not axioms) from ω-germ transformers via normalized sequential action on tails; (vi) characterisations of σ-anti-holomorphy and idempotent-supported holomorphy (factoring through exactly one lemniscate lobe); and (vii) construction of the holomorphic endomorphism category HolEnd_τ via pre-Yoneda collapse, with identity and composition proven as internal theorems. The scope is deliberately restricted to the *boundary algebra* (full Hartogs continuation, interior points, topology-geometry parallel readouts, and τ-Navier–Stokes are deferred to later work). The development is entirely boundary-first and diagonal-free; it imports only finite witness predicates, stabilisation, and the Hinge-4 scalar algebra $\mathbb{D}$. Lean-formalisation is planned in `TauLib.BookII.Holomorphy`.

Keywords split-complex holomorphy, ω-germ transformers, wave-equation Cauchy–Riemann, boundary algebra, diagonal discipline, earned category theory, pre-categorical kernel, no Cartesian product, pre-Yoneda collapse, hyperbolic number plane, Panta Rhei hinge paper

MSC 2020 Mathematics Subject Classification: 32A10, 30G35, 18A05, 03F65, 18B25

CONTENTS

1. INTRODUCTION AND STATEMENT OF RESULTS

1.1 Position in the Panta Rhei hinge-paper bundle

This paper is **Hinge 5** of the eight-paper Panta Rhei foundational bundle accompanying the 2nd Edition of the series [?, ?, ?]. 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* [?] — unique tower-atom decomposition $X = (A \uparrow\uparrow C)^B \cdot D$ in ZFC and Category τ; supplies the ABCD coordinate atoms.

Hinge 2: *Prime Polarity* [?] — classifies primes into B/C channels via the Legendre symbol $(2/p)$ mod 8.

Hinge 3: *Master Constant ι_τ* [?] — derives $\iota_\tau = 2/(\pi + e) \approx 0.341304$ as the unique σ-fixed crossing-germ scalar on the lemniscate.

Hinge 4: *The Split-Complex Boundary Algebra* [?] — establishes $\mathbb{D} = \mathcal{R}'_\partial[j]/(j^2-1)$ as the unique τ-admissible scalar algebra and the common algebraic home of Hinges 1–3.

Hinge 5: *τ-Holomorphy on the Boundary Algebra (this paper)* — installs τ-holomorphy as the ontological primary (before any notion of mapping), proves the wave-equation Cauchy–Riemann theorem, derives the earned categorical apparatus, and constructs HolEnd_τ.

Hinge 6: *The τ-Topos and Its Four-Valued Internal Logic* [?] — builds the internal logic Truth4 over the split-complex idempotent sublattice $B_\sigma(\mathbb{D})$ and internalises it as the subobject classifier Ω_τ of the countable τ-topos.

Hinge 7: *Address Resolution, Not Calculation* [?] — canonical-address NF confluence (Church–Rosser for the τ-kernel), genealogical DAG, Cayley word metric, ontic ultrametric; discharges the "modulo Hinge 7" caveats carried by the present paper's pre-Yoneda collapse theorem.

Hinge 8: *The τ-Kernel as Foundational Architecture* [?] — 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.

Each paper is standalone-readable, but the recommended sequencing above is the natural dependency order. The present paper imports from Hinge 4 the unique boundary algebra $\mathbb{D}$ and the canonical σ-involution; it supplies to Hinge 6 the earned categorical machine on which τ-topos morphisms are then built.

1.2 Boundary-first, pre-categorical holomorphy

Classical analysis proceeds

$$\text{Sets} \to \text{Maps} \to \text{Continuous Maps} \to \text{Differentiable Maps} \to \text{Holomorphic Maps}.$$

Each step imports machinery from the previous: Cartesian product, function-graph-as-slice, metric, derivative. The present paper *inverts this dependency order*. We define τ-holomorphy directly, without any prior notion of "mapping," "tuple," or "Cartesian product," as a certified transformer on the ω-tails of the boundary algebra. From this primary definition the entire categorical apparatus (composition, identity, associativity, functoriality) and the signature PDE of holomorphy (the Cauchy–Riemann equations, here realised as the hyperbolic *wave equation*) emerge as theorems.

Three inversions fuse into one structural move:

(i) *Holomorphy before mappings* — ω-germ transformers precede the notion "function";

(ii) *No Cartesian product / no graph* — the diagonal discipline (no free function-graph-as-slice through $D \times C$) is *precisely what protects* the split-complex idempotent structure of $\mathbb{D}$ from collapsing to elliptic $\mathbb{C}$;

(iii) *Wave-equation Cauchy–Riemann* — the split-complex CR decouples into the hyperbolic wave equation $\partial_t^2 f = \partial_x^2 f$, not the elliptic Laplace equation $\partial_t^2 f + \partial_x^2 f = 0$.

These three are *the same move viewed from three angles*: avoid the unearned diagonal, land in $\mathbb{D}$ (Hinge 4's forced algebra), and read out wave-type instead of Laplace-type PDEs.

1.3 What this paper is — and what it is not

The scope is deliberately restricted to **the boundary algebra itself**. We do *not* develop:

- Full Hartogs continuation and interior-point construction — these are deferred to a Book-II paper, where boundary germs get extended to interior sections.
- Topology-geometry parallel readouts (why the ultrametric and Euclidean geometry are both induced by holomorphy) — Book II.
- τ-Navier–Stokes as local Hartogs continuation — Book III.
- Riemann-Mapping-Theorem-analogue for lemniscate boundary automorphisms — Book II.

The present paper earns the *ontological primary* and its immediate categorical consequences; extensions are handed off cleanly to later foundational work.

1.4 Main theorems (summary)

Theorem 1.1 (Admissible boundary germs form a certified class **[τ-Effective]**). *For each carrier $X \in \mathrm{Obj}(\tau)$ there is a type $\mathsf{Germ}_\tau(X)$ of* admissible boundary germs *represented by NF-coded tail transformers $c\colon \Omega_{\mathrm{tail}} \to \Omega_{\mathrm{tail}}$ satisfying a tail-stability predicate* Stable *and a tail-independence predicate beyond a finite witness depth. Germ extensional equality is decidable-by-witness and compatible with stabilisation.*

Theorem 1.2 (Holomorphy before mappings **[τ-Effective]**). *τ-holomorphic maps are precisely the admissible tail transformers:*

$$\mathsf{Hol}_\tau(X,Y) \;:=\; \{c \in \mathsf{Code} \mid \mathsf{Typed}(X,Y,c) \wedge \mathsf{Stable}(X,Y,c) \wedge \textit{tail-independent}\}.$$

No function space $\mathrm{Hom}(X,Y)$, no Cartesian product $X \times Y$, and no function-graph are required. Interior "functions" are downstream shadows of the transformers on $\sim$-classes. The collection $\mathsf{Hol}_\tau(X,Y)$ is countable under the witness discipline.

Theorem 1.3 (Earned scalar codomain: $\mathbb{D}$ is forced **[τ-Effective]**). *Among all associative $\mathcal{R}'_\partial$-algebras S compatible with the ω-germ transformer axioms and the σ-involution inherited from the lemniscate, the unique such S (up to canonical isomorphism) is the split-complex boundary algebra $\mathbb{D} = \mathcal{R}'_\partial[j]/(j^2-1)$ of Hinge 4 [**?**]. Elliptic alternatives ($\mathbb{C}$, quaternionic slices) fail by the idempotent-obstruction argument of Hinge 4 Theorem 1.7.*

Theorem 1.4 (Wave-equation Cauchy–Riemann **[τ-Effective]**). *For a τ-holomorphic map $f \in \mathsf{Hol}_\tau(X,\mathbb{D})$ written in the basis $f(z) = u(z) + j\,v(z)$ with $u,v\colon X \to \mathcal{R}'_\partial$ and $j^2 = +1$, the split-complex Cauchy–Riemann equations*

$$\partial_t u \;=\; \partial_x v, \qquad \partial_x u \;=\; \partial_t v$$

hold, and each component u, v satisfies the hyperbolic wave equation

$$\partial_t^2\, u \;=\; \partial_x^2\, u, \qquad \partial_t^2\, v \;=\; \partial_x^2\, v.$$

This is a structural consequence of $j^2 = +1$ — the signature inherited from Hinge 4 — rather than of any imposed regularity.

Theorem 1.5 (Diagonal discipline: the no-Cartesian-product principle **[τ-Effective]**). *Within the ω-germ transformer framework, there is no admissible construction that produces a "function graph" as a slice through a Cartesian product $X \times Y$ of carriers. Consequently, the idempotent collapse that would force the boundary algebra toward elliptic $\mathbb{C}$ (Hinge 4 Theorem 1.7 alternative) is structurally precluded. The diagonal-free discipline is the exact structural obstruction that protects the split-complex signature $j^2 = +1$.*

Theorem 1.6 (The earned categorical machine **[τ-Effective]**). *Using ω-germ transformers as the only primitive, the following categorical structure is derived as* theorems*:*

(a) ***Composition.*** *The composite of two admissible transformers is again admissible, and composition is given by normalized sequential action on tails with code concatenation.*

(b) ***Identity.*** *The identity transformer $\mathrm{id}_X \in \mathsf{Hol}_\tau(X,X)$ is the canonical tail-fixing NF code.*

(c) ***Associativity.*** *Associativity of composition holds up to canonical NF reduction and confluence of the rewriting system on codes.*

(d) ***Functoriality.*** *The assignment $X \mapsto \mathsf{Hol}_\tau(X,-)$ is a functor on the probe category of carriers, with boundary-preserving closure as its naturality condition.*

No category axioms are imposed; they are earned.

Theorem 1.7 (σ-anti-holomorphy and idempotent-supported holomorphy **[τ-Effective]**). *Anti-holomorphic transformers in τ are σ-conjugates of holomorphic ones: $\bar{f} := \sigma \circ f \circ \sigma \in \mathsf{Hol}_\tau(X,Y)$ for every $f \in \mathsf{Hol}_\tau(X,Y)$. Moreover, every admissible τ-holomorphic map into $\mathbb{D}$ factors uniquely through exactly one of the two idempotent lobes $e_+\mathbb{D}$ or $e_-\mathbb{D}$:*

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

This is the idempotent-supported-holomorphy theorem, inherited from Hinge 4's four-atom dictionary and lifted to the level of transformers.

Theorem 1.8 (HolEnd$_\tau$ via pre-Yoneda collapse, modulo Hinge 7 canonical-address NF confluence **[τ-Effective]**). *The holomorphic endomorphism category*

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

is constructible entirely from ω-germ transformers and a canonical pre-Yoneda collapse (embedding Hol_τ *into the boundary addressable objects of Hinge 1). Identity, composition, endomorphisms, and automorphisms are theorems of the construction, not axioms. Under the σ-equivariance refinement,* $\mathsf{HolEnd}_\tau^\sigma$ *(the σ-fixed endomorphism monoid) is the natural host for the universality statements of Hinge 3's master constant. The "pre-Yoneda collapse" used here is stated modulo the canonical-address normalisation to be established in Hinge 7; see Theorem* **??** *(§**??**) for the precise form and the Hinge-7 dependency.*

1.5 Hinge-integration theorem

Theorem 1.9 (Hinge 5 integration tabulation **[τ-Effective]**). *Hinge 5 is positioned in the Panta Rhei hinge bundle by the following four backward dependencies and two forward obligations, each of which is made formally precise in the sections indicated.*

- ***From Hinge 4*** *(§§**??**, **??**): the unique boundary algebra $\mathbb{D}$, its four idempotent atoms $\{0, e_+, e_-, 1\}$, the σ-involution, and the uniqueness / elliptic-exclusion theorems are imported as starting data, not rederived.*
- ***To Hinge 6*** *(§§**??**, **??**): the earned categorical machine (Theorem **??**) is the base on which the τ-topos is built; the idempotent-supported holomorphy theorem (Theorem **??**) provides the subobject classifier's four truth values $\{0, e_+, e_-, 1\}$.*
- ***To Hinge 7*** *(§**??**): the pre-Yoneda collapse (Theorem **??**) links* HolEnd_τ *to the canonical-address NF confluence of Hinge 7; transformers are addressed, not generated.*
- ***Backward to Hinges 1–3*** *(Theorem **??**): the ABCD coordinates [**?**], the prime-polarity character [**?**], and the master constant ι_τ [**?**] all appear as specific admissible transformers in* Hol_τ *or as σ-fixed scalars in the codomain $\mathbb{D}$.*

The present paper completes the bundle's fifth *hinge: the earned categorical home in which Hinges 1–4 are internalised as transformers, and from which Hinges 6–7 inherit their logical and address-theoretic superstructures.*

1.6 Lean roadmap (preview)

Full formalisation is targeted at `TauLib.BookII.Holomorphy` [**?**] in the Lean 4 proof assistant [**?**], comprising the following planned modules:

- `Tails.lean` — ω-tails, prefixes, tail-equivalence predicates.
- `Germs.lean` — admissible boundary germ type Germ_τ.
- `HolMaps.lean` — $\mathsf{Hol}_\tau(X, Y)$ as certified transformer type.
- `EarnedScalar.lean` — the scalar-codomain uniqueness (forcing $\mathbb{D}$).
- `WaveCR.lean` — split-complex Cauchy–Riemann equations and wave-equation theorem.
- `DiagonalDiscipline.lean` — no-Cartesian-product obstruction.
- `EarnedCat.lean` — earned composition, identity, associativity, functoriality.
- `HolEnd.lean` — HolEnd_τ and $\mathsf{HolEnd}_\tau^\sigma$ via pre-Yoneda collapse.

See §**??** for the full proof-chain sketch.

2. PRELIMINARIES: TAILS, PREFIXES, AND TAIL-EQUIVALENCE

Definition 2.1 (ω-tails). *An ω-*tail *is an infinite coherent prefix chain over the τ-native token alphabet. We denote the type of tails by Ω_{tail}. For each $k \in \text{Idx}$ and prefix token σ of length k we write $\mathsf{Pref}_{k,\sigma}(t)$ for the decidable predicate "t has prefix σ at depth k".*

Definition 2.2 (Prefix agreement and tail-equivalence). *For tails $t, t' \in \Omega_{\text{tail}}$ define:*

$$t \equiv_k t' \iff t \text{ and } t' \text{ agree on all prefixes up to depth } k,$$

and the full tail-equivalence

$$t \sim t' \iff \forall k \in \text{Idx},\ t \equiv_k t'.$$

Lemma 2.3 (Tail equivalence is an equivalence relation **[Established]**). *$\sim$ is reflexive, symmetric, and transitive.*

Lean-grade sketch. Immediate from the corresponding properties of prefix agreement. □

Definition 2.4 (Carrier compatibility). *For each $X \in \mathrm{Obj}(\tau)$ there is a predicate $\mathsf{Tail}_X : \Omega_{\text{tail}} \to \mathsf{Prop}$, invariant under $\sim$, interpreted as "t is an admissible tail for carrier X".*

Remark 2.5 (What we do *not* assume). We emphasise the minimality of the setup:

- **No Cartesian products.** The symbol "$X \times Y$" is undefined; pairs are typed orbit slots, not tuples.
- **No function spaces.** $\mathrm{Hom}(X, Y)$ as an object does not exist at this layer; admissible transformers will play its role intensionally but never as an exponential object.
- **No set comprehension.** Collections are either finite enumerations (witness-coded) or earned as $\sim$-quotients of already-admissible tails.
- **No real-analytic substrate.** All infinitary content is stabilisation of finite witnesses.

These are the four prohibitions that together *protect* the split-complex structure of the boundary algebra from elliptic collapse (Theorem **??**, proven in §**??**).

3. THE ONTOLOGICAL PRIMARY: τ-HOLOMORPHIC MAPS AS CERTIFIED TRANSFORMERS

Definition 3.1 (Tail transformer code). *A* tail transformer *is given by a code $c \in \mathsf{Code}$ with semantics $[\![c]\!] : \Omega_{\text{tail}} \to \Omega_{\text{tail}}$.*

Definition 3.2 (Typing). *A code c is* typed *$X \to Y$ if*

$$\mathsf{Typed}(X, Y, c) \; :\Longleftrightarrow \; \forall t \in \Omega_{\text{tail}},\; \mathsf{Tail}_X(t) \Rightarrow \mathsf{Tail}_Y([\![c]\!](t)).$$

Definition 3.3 (Tail stability). *A code c is* tail-stable *on $X \to Y$ if*

$$\mathsf{Stable}(X, Y, c) \; :\Longleftrightarrow \; \forall t, t' \in \Omega_{\text{tail}},\; \mathsf{Tail}_X(t) \wedge \mathsf{Tail}_X(t') \wedge (t \sim t') \Rightarrow ([\![c]\!](t) \sim [\![c]\!](t')).$$

Definition 3.4 (Tail-independence). *A code c is* tail-independent beyond depth k_0 *on $X \to Y$ if there exists $k_0 \in \mathrm{Idx}$ such that whenever $\mathsf{Tail}_X(t)$, $\mathsf{Tail}_X(t')$ and $t \equiv_{k_0} t'$, then $[\![c]\!](t) \sim [\![c]\!](t')$.*

Definition 3.5 (τ-holomorphic map). *A* τ-holomorphic map *from carrier X to carrier Y is a code $c \in \mathsf{Code}$ satisfying $\mathsf{Typed}(X, Y, c)$, $\mathsf{Stable}(X, Y, c)$, and tail-independence beyond some finite depth. The type of τ-holomorphic maps is denoted*

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

Remark 3.6 (Why this deserves the name "holomorphic"). A τ-holomorphic map as just defined is not presented as a rule from an ambient set to another; it is an intensional coherence certificate on ω-tails. Four structural features justify the name "holomorphic" despite the absence of derivatives, analytic charts, or function spaces:

(i) *Finite-depth determinacy.* By tail-independence, the ω-level behaviour of f is determined by its behaviour at some finite witness-depth k_0. This is the τ-native analogue of the *identity theorem* for holomorphic functions (determinacy by a small amount of data).

(ii) *Rigidity under stabilisation.* Stable enforces that equivalent tails go to equivalent tails; f respects the full $\sim$-structure. This is the τ-native analogue of *coherence along boundary data* (the Hartogs-type boundary-to-interior principle, restricted here to boundary-to-boundary).

(iii) *No unearned analytic content.* No derivative, no power series, no $\bar{\partial}$-operator is imposed; the Cauchy–Riemann structure is *derived* later (Theorem **??**) from the codomain algebra.

(iv) *Categorical closure under composition.* The admissible transformers are closed under composition (§**??**), giving a category whose name — HolEnd_τ — is again traditional.

Collectively, these features give the τ-native object every structural virtue classical holomorphy is prized for (rigidity, boundary-determinacy, algebraic closure under composition) without importing any of the baggage that classical holomorphy carries (Cartesian products, metric substrates, analytic charts). The name is earned, not imposed.

Remark 3.7 (Countability under the witness discipline). $\mathsf{Hol}_\tau(X, Y)$ is countable: each element is represented by a finite NF-coded program c, and the admissibility predicates (Typed, Stable, tail-independence) are decidable by finite witness examination. No cardinality jumps occur; in particular, there is no uncountable function-space.

Remark 3.8 (Why tail transformers are not smuggled function-spaces). An obvious objection deserves a direct response: a tail transformer code c has semantics $[\![c]\!]: \Omega_{\text{tail}} \to \Omega_{\text{tail}}$, which *looks like* a function from one set to another. If we have banished function spaces, how is this admissible? The answer turns on the distinction between *set-theoretic function* and *intensional code semantics*:

(a) A set-theoretic function is a graph $\{(t, f(t)) : t \in \Omega_{\text{tail}}\} \subseteq \Omega_{\text{tail}} \times \Omega_{\text{tail}}$; building it requires the Cartesian product $\Omega_{\text{tail}} \times \Omega_{\text{tail}}$ as a carrier, which (DD1)–(DD2) of the diagonal discipline (§??) forbid.

(b) An NF-coded tail transformer is a finite code $c \in \mathsf{Code}$ together with a primitive-recursive evaluator that produces an output tail from each input tail *by stepwise token substitution*. No ambient product is formed; the evaluator is a meta-level rewriting procedure, not a set of ordered pairs.

The distinction is the same one that linear logic makes between *types* (intensional, no free contraction) and *sets* (extensional, free diagonal access). The $\Omega_{\text{tail}} \to \Omega_{\text{tail}}$ arrow in code semantics is a type-theoretic arrow, not a set-theoretic one; it does not bring back the Cartesian-product structure that (DD1)–(DD2) rule out. This type–set distinction is the same one that constructive and type-theoretic foundations have long made explicit: see [?, ?] for the constructive-mathematics tradition, and [?] for its modern univalent/homotopy-type-theoretic formulation. The critical test is whether quantification "$\forall t \in \Omega_{\text{tail}}$" in Definitions ??–?? can be internalised as membership in a set — and it cannot: Ω_{tail} is given to us as a *type* of admissible coherence certificates, and the quantification is genuinely parametric over such certificates, not existential over a comprehension-defined collection. The τ-kernel's countability and type-theoretic discipline (§?? Remark ??) guarantee this distinction is maintained throughout.

4. EARNED SCALAR CODOMAIN: $\mathbb{D}$ IS FORCED

The holomorphic map definition of §?? (Definition ??) takes values in a scalar algebra S that was, at that point, a parameter of the construction. We now discharge this parameter. In this section we prove that among all $\mathcal{R}'_\partial$-algebras compatible with the ω-germ transformer axioms of §3 and with the σ-structure inherited from the lemniscate $\mathbb{L}$, there is exactly one such algebra (up to canonical isomorphism): the split-complex boundary algebra

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

of Hinge 4 [?]. The scalar codomain of τ-holomorphy is not a modelling choice; it is forced.

4.1 What "compatible scalar codomain" means

Before we can state a uniqueness theorem, we must isolate which features of S are essential to the transformer framework and which are incidental. The holomorphic map definition Definition ?? requires that the stabilised tail of each ω-germ transformer have a well-defined limit value, that this value participate in the σ-involution inherited from the lobe-swap of $\mathbb{L}$, and that it decompose along the B/C polarity channels of Hinge 2 [?]. We formalise these requirements as follows.

Definition 4.1 (Compatible scalar codomain, [τ-Effective]). *V.D.ESC.1 A* compatible scalar codomain *for ω-germ transformers over the boundary ring $\mathcal{R}'_\partial$ is a pair (S, σ_S) consisting of an $\mathcal{R}'_\partial$-algebra S and an $\mathcal{R}'_\partial$-linear map $\sigma_S : S \to S$, satisfying the following four axioms.*

(CC1) Base ring. *S is a commutative, associative, unital $\mathcal{R}'_\partial$-algebra. The structure map $\mathcal{R}'_\partial \hookrightarrow S$ is a unital ring homomorphism, and S is finite-rank and torsion-free as an $\mathcal{R}'_\partial$-module.*

(CC2) Canonical involution. *σ_S is an $\mathcal{R}'_\partial$-algebra involution, i.e. $\sigma_S^2 = \mathrm{id}_S$, $\sigma_S(1) = 1$, and $\sigma_S(xy) = \sigma_S(x)\sigma_S(y)$. Moreover σ_S is the algebraic shadow of the lemniscate lobe-swap $\sigma_{\mathbb{L}}$, in the sense that the transformer space $\mathsf{Hol}_\tau(X, S)$ is closed under post-composition with σ_S, and this post-composition corresponds via the ω-germ evaluation to the lobe-swap action on $\mathbb{L}$ [?].*

(CC3) Stable tail admissibility. *For every ω-germ transformer $c \in \mathsf{Hol}_\tau(X, S)$, the stabilised tail $[\![c]\!]$ lands in S and is $\sim$-stable: if $c \sim c'$ then $[\![c]\!] = [\![c']\!]$. Equivalently, S hosts the equivalence-class values of* Stable ∘ Tail *without collapse.*

(CC4) Bipolar idempotent decomposition. *There exists a pair $e_+, e_- \in S$ of orthogonal idempotents with $e_+ + e_- = 1$ and $e_+ \cdot e_- = 0$, such that the σ_S-action swaps them: $\sigma_S(e_+) = e_-$. The decomposition $S = e_+S \oplus e_-S$ is the algebraic incarnation of the B/C-channel bipartition of [?].*

Remark 4.2 (Why these four and not more). The four axioms (CC1)–(CC4) are not ad hoc. (CC1) fixes the base over which scalars live; without it the transformer values would not be $\mathcal{R}'_\partial$-compatible with the ambient boundary ring. (CC2) is the condition for σ-symmetry of the transformer space, forced by the lemniscate's lobe-swap [?]. (CC3) is the admissibility condition built into Definition ?? (items Stable and $\sim$ of §3). (CC4) is the algebraic consequence of the B/C bipartition from prime polarity [?], which every τ-level structure must respect.

Any additional axiom — completeness, norm, order, differentiable structure — would exceed what the transformer framework demands and would no longer track the true degrees of freedom. Conversely, dropping any one of (CC1)–(CC4) breaks either the base compatibility, the σ-symmetry, the tail stability, or the polarity decomposition. Thus (CC1)–(CC4) are minimal and sufficient.

Remark 4.3 (Isomorphism of compatible codomains). We say two compatible codomains (S, σ_S) and $(S', \sigma_{S'})$ are *canonically isomorphic over* σ if there is an $\mathcal{R}'_\partial$-algebra isomorphism $\varphi : S \to S'$ with $\varphi \circ \sigma_S = \sigma_{S'} \circ \varphi$ and $\varphi(e_+^S) = e_+^{S'}$, $\varphi(e_-^S) = e_-^{S'}$. Uniqueness is always stated up to this notion of isomorphism.

4.2 The earned scalar codomain theorem

The four axioms (CC1)–(CC4) admit exactly one solution. We now state and prove this.

Theorem 4.4 (Earned Scalar Codomain, **[τ-Effective]**). *V.T.ESC.MAIN Among all pairs (S, σ_S) satisfying the compatibility axioms* (CC1)–(CC4) *of Definition* ??*, there exists a unique such pair (up to canonical isomorphism of $\mathcal{R}'_\partial$-algebras over σ), namely the split-complex boundary algebra*

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

Explicitly: there is a canonical $\mathcal{R}'_\partial$-algebra isomorphism $\varphi : S \xrightarrow{\sim} \mathbb{D}$ with $\varphi \circ \sigma_S = \sigma_{\mathbb{D}} \circ \varphi$, sending $e_+^S \mapsto e_+ = (1+j)/2$ and $e_-^S \mapsto e_- = (1-j)/2$.

Proof. The proof proceeds in three steps. Step 1 reduces (CC1)–(CC4) to the boundary-algebra axioms (C1)–(C4) of Hinge 4. Step 2 invokes Hinge 4's canonical uniqueness theorem. Step 3 verifies (CC3) on the candidate $\mathbb{D}$.

Step 1: (CC1)+(CC2)+(CC4) imply the Hinge 4 boundary-algebra axioms.

Recall [?, Def. 1.X, Axioms (C1)–(C4)]: a *boundary algebra* is a commutative associative unital $\mathcal{R}'_\partial$-algebra A equipped with an $\mathcal{R}'_\partial$-algebra involution σ_A and a pair of orthogonal idempotents e_+^A, e_-^A with $e_+^A + e_-^A = 1$, $\sigma_A(e_+^A) = e_-^A$, and $A = e_+^A A \oplus e_-^A A$ with each summand a free $\mathcal{R}'_\partial$-module of rank 1.

We verify each of Hinge 4's (C1)–(C4) for (S, σ_S):

- (C1) S is a commutative associative unital $\mathcal{R}'_\partial$-algebra: given by (CC1).
- (C2) σ_S is an $\mathcal{R}'_\partial$-algebra involution: given by (CC2).
- (C3) Idempotents e_+^S, e_-^S with $e_+^S + e_-^S = 1$, orthogonal, and swapped by σ_S: given by (CC4).
- (C4) Rank-1 summands: from (CC1), S is finite-rank and torsion-free over $\mathcal{R}'_\partial$. The decomposition $S = e_+^S S \oplus e_-^S S$ from (CC4) is therefore a direct sum of $\mathcal{R}'_\partial$-submodules. Because σ_S exchanges the two summands isomorphically (being an involution that swaps the idempotents), the two summands have equal $\mathcal{R}'_\partial$-rank. If that common rank were ≥ 2, each summand would contain a nontrivial idempotent projection refining e_+^S or e_-^S; but (CC4) fixes the idempotents uniquely as the σ_S-swap pair, so no finer refinement exists. Hence each summand has rank 1, and (C4) holds.

Step 2: Canonical uniqueness.

By [?, Thm. 1.6, Canonical Uniqueness of $\mathbb{D}$], any boundary algebra $(A, \sigma_A, e_+^A, e_-^A)$ satisfying (C1)–(C4) is canonically isomorphic, as an $\mathcal{R}'_\partial$-algebra with involution and distinguished idempotents, to $(\mathbb{D}, \sigma_{\mathbb{D}}, e_+, e_-)$. Applying this to $(S, \sigma_S, e_+^S, e_-^S)$ gives the required canonical isomorphism $\varphi : S \xrightarrow{\sim} \mathbb{D}$.

Step 3: (CC3) holds for $\mathbb{D}$.

It remains to check that $\mathbb{D}$ itself satisfies the stable-tail admissibility axiom (CC3), confirming that $\mathbb{D}$ is in fact a compatible scalar codomain (and not merely the unique candidate under (CC1), (CC2), (CC4)).

Let $c \in \mathsf{Hol}_\tau(X, \mathbb{D})$ be an ω-germ transformer with values in $\mathbb{D}$. Decompose

$$c(x) \;=\; e_+\, c_+(x) \;+\; e_-\, c_-(x), \qquad c_\pm(x) \in \mathcal{R}'_\partial,$$

using the idempotent decomposition $\mathbb{D} = e_+\mathbb{D} \oplus e_-\mathbb{D} \cong \mathcal{R}'_\partial \oplus \mathcal{R}'_\partial$. Because e_+ and e_- are central idempotents and the tail functor Tail is $\mathcal{R}'_\partial$-linear, we have

$$\mathsf{Tail}(c) \;=\; e_+\,\mathsf{Tail}(c_+) \;+\; e_-\,\mathsf{Tail}(c_-).$$

Stabilisation Stable commutes with scalar multiplication by e_+, e_- (the idempotent structure passes through the $\sim$-equivalence because orthogonal idempotents are preserved by any $\mathcal{R}'_\partial$-algebra endomorphism). Therefore

$$[\![c]\!] \;=\; e_+\,[\![c_+]\!] \;+\; e_-\,[\![c_-]\!] \;\in\; \mathbb{D}.$$

Both $[\![c_+]\!]$ and $[\![c_-]\!]$ are $\sim$-stable in $\mathcal{R}'_\partial$ by the one-dimensional case of Definition **??**. Consequently $[\![c]\!]$ is $\sim$-stable in $\mathbb{D}$.

Combining Steps 1–3: every (S, σ_S) satisfying (CC1)–(CC4) is canonically isomorphic to $\mathbb{D}$, and $\mathbb{D}$ itself satisfies (CC1)–(CC4). Uniqueness and existence together yield the theorem. □

Remark 4.5 (Role of dyadic localisation)**.** The localisation $\mathcal{R}'_\partial = \mathcal{R}_\partial[1/2]$ is used in Step 2 when writing the canonical idempotents $(1 \pm j)/2$. Without inverting 2, the boundary ring $\mathcal{R}_\partial$ cannot host the e_+/e_- decomposition of (CC4). This is why Hinge 4 localises to $\mathcal{R}'_\partial$ at the outset — the localisation is forced by (CC4) in exactly the same way that (CC4) itself is forced by B/C-polarity. See [**?**, §2] for the detailed discussion.

4.3 Elliptic exclusion for transformers

Theorem **??** selects $\mathbb{D}$ as the unique scalar codomain. But the classical impulse, encountering a two-dimensional $\mathcal{R}'_\partial$-algebra with involution, is to reach instead for the elliptic algebra

$$\mathbb{C}_\partial \;:=\; \mathcal{R}'_\partial[i]/(i^2+1), \qquad \sigma_{\mathbb{C}_\partial}(i) = -i.$$

We now show directly, in the transformer language, that $\mathbb{C}_\partial$ fails (CC4) and hence fails compatibility.

Proposition 4.6 (Elliptic Exclusion in Transformer Language, **[Established]**)**.** V.P.ESC.ELL *The elliptic algebra $\mathbb{C}_\partial = \mathcal{R}'_\partial[i]/(i^2+1)$, equipped with the complex-conjugation involution $i \mapsto -i$, does not satisfy axiom* (CC4) *of Definition* **??**. *In particular, $\mathbb{C}_\partial$ is not a compatible scalar codomain for ω-germ transformers.*

Proof. The naive candidates for a σ-swap idempotent pair in $\mathbb{C}_\partial$ are

$$f_+ \;=\; \frac{1+i}{2}, \qquad f_- \;=\; \frac{1-i}{2}.$$

These are σ-swapped and sum to 1. However they are not idempotent. Compute:

$$f_+^2 \;=\; \left(\frac{1+i}{2}\right)^2 \;=\; \frac{1+2i+i^2}{4} \;=\; \frac{1+2i-1}{4} \;=\; \frac{i}{2} \;\neq\; \frac{1+i}{2} \;=\; f_+.$$

Symmetrically $f_-^2 = -i/2 \neq f_-$. Moreover, any idempotent $e \in \mathbb{C}_\partial$ with $e + \sigma(e) = 1$ and $e \cdot \sigma(e) = 0$ must satisfy $e(1-e) = 0$ and $e^2 = e$, forcing $e \in \{0, 1\}$ because $\mathbb{C}_\partial$, viewed as an algebra over the field of fractions of $\mathcal{R}'_\partial$, is the field of complex-like numbers and hence contains no nontrivial idempotents. The only idempotents are the trivial 0 and 1, and these are not σ-swapped. Therefore no pair $(e_+^{\mathbb{C}_\partial}, e_-^{\mathbb{C}_\partial})$ satisfying (CC4) exists in $\mathbb{C}_\partial$. □

Remark 4.7 (What goes wrong geometrically). Proposition **??** has a clean transformer interpretation. If one attempted to build ω-germ transformers into $\mathbb{C}_\partial$, the σ-action would act as rotation $i \mapsto -i$ rather than as a genuine swap of orthogonal sectors. There is no B-sector and C-sector to carry the prime-polarity bipartition; the would-be sectors fuse into a single rotational orbit. Hinge 2's B/C bipartition [**?**] is structurally incompatible with rotational symmetry: B and C are orthogonal, not complex-conjugate. The split-complex $\mathbb{D}$ carries a reflective σ (across the j-line) that genuinely exchanges two half-lines, while the elliptic $\mathbb{C}_\partial$ carries a rotational σ that fixes only two points. Idempotent-supported holomorphy (§**??**) requires genuine sector decomposition; it cannot survive the elliptic collapse.

Corollary 4.8 (No rotational holomorphy, **[Established]**). *There is no ω-germ transformer framework* $\mathsf{Hol}_\tau(X, \mathbb{C}_\partial)$ *compatible with the* B/C*-bipartition of* [**?**]*. Any attempt to define such a framework either violates (CC4) or collapses the two sectors into a single rotational orbit, in which case the Hinge 2 polarity structure is lost.*

Proof. Immediate from Proposition **??** and Remark **??**. □

The exclusion of $\mathbb{C}_\partial$ is the algebraic mirror of Hinge 4's geometric elliptic exclusion [**?**, Thm. 1.7]. The two exclusions are facets of the same underlying fact: the lemniscate $\mathbb{L} = S^1 \vee S^1$ has a reflective σ, not a rotational one, and scalar codomains of τ-holomorphy must track this reflective structure algebraically. For the classical split-complex literature outside the τ-context, see [**?**, **?**, **?**].

4.4 Interpretation: the structural moral

Theorem **??** closes a loop that begins in Book I with the seven axioms K0–K6 and the five generators, passes through Book II's definition of the lemniscate, and arrives here at the uniquely determined scalar codomain.

Remark 4.9 (The scalar codomain is *earned*, not chosen). One might have expected the scalar codomain of τ-holomorphy to be a modelling choice — $\mathbb{C}$, $\mathbb{R}$, $\mathbb{D}$, $\mathbb{H}$, or some other two-or-four-dimensional algebra — to be selected for convenience or aesthetic preference. Theorem **??** proves that this expectation is wrong. The scalar codomain is determined by:

1. The base ring $\mathcal{R}'_\partial$ (fixed by the boundary-ring construction of Hinge 4 [**?**]).
2. The σ-involution (fixed by the lemniscate lobe-swap of Hinge 3 [**?**]).
3. The B/C-bipartition (fixed by prime polarity in Hinge 2 [**?**]).
4. The tail-admissibility (fixed by ω-germ stability in §3 [**?**]).

Each of these four inputs is a τ-structural commitment made elsewhere in the framework, not a choice made here. Their conjunction forces $\mathbb{D}$ exactly. The scalar codomain is earned by the preceding hinges, not imposed ab initio.

Remark 4.10 (Idempotents as features). A recurring theme across Hinges 2–4 is that the non-integral elements of $\mathbb{D}$ — the zero-divisors e_+ and e_- — are not defects but features. The classical complex-analytic reflex is to prize fields over rings with zero-divisors; in the τ-framework this reflex is reversed. The idempotents e_+, e_- *are* the B and C sectors; their orthogonality *is* the polarity bipartition; their exchange under $\sigma_{\mathbb{D}}$ *is* the lobe-swap. A field-valued codomain would destroy precisely what the framework is built to see. As [**?**, §1] puts it: "the zero-divisors are where the geometry happens."

Theorem **??** reinforces this: insisting on a field forces $\mathbb{C}_\partial$, which fails (CC4), which breaks the sector decomposition, which erases polarity. The idempotents are not tolerated; they are required.

Remark 4.11 (Relation to the literature on hyperbolic numbers). The algebra $\mathbb{D}$ is known classically as the hyperbolic numbers, split-complex numbers, or double numbers; see [**?**, **?**, **?**] for surveys. These works study $\mathbb{D}$ as a curiosity or as a model of Minkowski geometry. Theorem **??** recontextualises $\mathbb{D}$: it is not a curiosity but the *forced* scalar codomain for any τ-level holomorphy theory. What was marginal in classical analysis is central in τ-analysis, and what was central in classical analysis ($\mathbb{C}$, with its rotational σ) is excluded in τ-analysis by Corollary **??**.

4.5 Corollary: sector decomposition of holomorphic maps

As an immediate byproduct of Theorem **??**, every holomorphic map with values in $\mathbb{D}$ decomposes along the B/C sectors. This foreshadows the idempotent-supported-holomorphy theorem of §**??**.

Corollary 4.12 (Canonical sector decomposition, **[τ-Effective]**). *V.C.ESC.SECTOR Let X be any τ-object. Every holomorphic map $f \in \mathsf{Hol}_\tau(X, \mathbb{D})$ decomposes uniquely and canonically as*

$$f \;=\; e_+ f_+ \;+\; e_- f_-,$$

where $f_+, f_- \in \mathsf{Hol}_\tau(X, \mathcal{R}'_\partial)$ are $\mathcal{R}'_\partial$-valued holomorphic maps, given explicitly by

$$f_+(x) \;=\; \tilde{\chi}(e_+ \cdot f(x)), \qquad f_-(x) \;=\; \tilde{\chi}(e_- \cdot f(x)),$$

under the canonical projection $\tilde{\chi} : e_+\mathbb{D} \oplus e\_\mathbb{D} \to \mathcal{R}'_\partial \oplus \mathcal{R}'_\partial$ determined by $\tilde{\chi}(e_+ a) = (a, 0)$ and $\tilde{\chi}(e\_b) = (0, b)$. The action of $\sigma_\mathbb{D}$ on f exchanges the two summands: $\sigma_\mathbb{D} \circ f = e_+ f_- + e_- f_+$.

Proof. For each $x \in X$, the value $f(x) \in \mathbb{D}$ admits the unique decomposition $f(x) = e_+ a + e\_b$ with $a, b \in \mathcal{R}'_\partial$, by the direct-sum structure $\mathbb{D} = e_+\mathbb{D} \oplus e\_\mathbb{D} \cong \mathcal{R}'_\partial \oplus \mathcal{R}'_\partial$. Define $f_+(x) = a$ and $f\_(x) = b$. Both f_+ and $f\_$ are ω-germ transformers with $\mathcal{R}'_\partial$-values because the idempotent projections $f \mapsto e_+ f$ and $f \mapsto e\_ f$ are $\mathcal{R}'_\partial$-linear operations that commute with Tail and preserve Stable (by the argument in Step 3 of Theorem **??**). Thus $f_\pm \in \mathsf{Hol}_\tau(X, \mathcal{R}'_\partial)$.

Uniqueness is immediate from the uniqueness of the direct-sum decomposition of $\mathbb{D}$. The $\sigma_\mathbb{D}$-action follows from $\sigma_\mathbb{D}(e_+) = e\_$ and linearity: $\sigma_\mathbb{D}(f(x)) = \sigma_\mathbb{D}(e_+ a + e\_b) = e\_a + e_+ b = e_+ f\_(x) + e\_f_+(x)$. ☐

Remark 4.13 (Foreshadowing §**??**). Corollary **??** is the precursor of the idempotent-supported-holomorphy theorem of §**??**, where we will show that the two components f_+ and $f\_$ can be independently chosen without violating τ-holomorphy. The $B_\sigma(\mathbb{D})$-sublattice structure on the σ-fixed idempotent spectrum will play the central role there. For now, Corollary **??** establishes the decomposition itself as a structural consequence of Theorem **??**; the freedom to mix and match sectors is deferred to §**??**.

Remark 4.14 (Closing remark of §4). With Theorem **??** and Corollary **??** in hand, the scalar codomain question raised by Definition **??** is fully resolved. The codomain is $\mathbb{D}$; it is forced; and every holomorphic map into it decomposes canonically along the sectors. The remainder of Hinge 5 (§§5–9) develops the consequences of this decomposition: function-theoretic properties, σ-equivariance, and the transition from scalar-valued holomorphy to the idempotent-supported extension of §**??**. The foundational work — identifying the codomain and proving its uniqueness — is complete. See also [**?**, **?**] for the broader place of this result in the Panta Rhei framework.

5. THE WAVE-EQUATION CAUCHY–RIEMANN THEOREM

5.1 Setup: split-complex differential structure

From Hinge 4 [**?**] we inherit the boundary-transformer algebra

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

together with the orthogonal idempotent basis

$$e_+ \;=\; \tfrac{1}{2}(1 + j), \qquad e_- \;=\; \tfrac{1}{2}(1 - j), \qquad e_+ \cdot e_- = 0, \qquad e_+ + e_- = 1.$$

As an $\mathcal{R}'_\partial$-module we therefore have the canonical decomposition

$$\mathbb{D} \;\cong\; \mathcal{R}'_\partial \oplus \mathcal{R}'_\partial j,$$

and every element $w \in \mathbb{D}$ admits the unique "real/imaginary" split $w = a + jb$ with $a, b \in \mathcal{R}'_\partial$.

Split-complex coordinates.. On the boundary transformer we introduce the split-complex variable

$$z \;=\; t + j\,x, \qquad t, x \in \mathcal{R}'_\partial,$$

with conjugation $\bar{z} := t - j\,x$. Because $j^2 = +1$, the quadratic form induced by multiplication is

$$z\bar{z} \;=\; (t + jx)(t - jx) \;=\; t^2 - x^2, \tag{1}$$

with a minus sign. This is the Minkowski (signature $(1,1)$) product [**?**, **?**, **?**]; in particular $z\bar{z} = 0$ occurs on the null-cone $\{t = \pm x\}$ rather than only at the origin, and the algebra $\mathbb{D}$ is consequently not a field.

Partial derivatives.. For functions $h\colon X \to \mathcal{R}'_\partial$ of the two underlying transformer coordinates (t,x) we write $\partial_t h$ and $\partial_x h$ for the componentwise partials in the sense of the $\mathcal{R}'_\partial$-module structure on X. These are the discrete-combinatorial shadows of translation on the light-cone lattice; we do *not* import any analytic notion of limit from classical $\mathbb{R}$-analysis. The operators ∂_t and ∂_x commute, $\partial_t\partial_x = \partial_x\partial_t$, by construction of the transformer-coordinate action.

A function $f\colon X \to \mathbb{D}$ is written uniformly as

$$f(z) \;=\; u(z) + j\,v(z) \qquad \text{with } u, v\colon X \to \mathcal{R}'_\partial.$$

In light of the idempotent basis we also record the equivalent form $f = f_+\,e_+ + f_-\,e_-$ where $f_+ = u + v$ and $f_- = u - v$; these are the *characteristic* or *lightcone* components.

5.2 Split-complex differentiability at the transformer level

We now specify what "$\mathbb{D}$-differentiable" means at the boundary transformer — without borrowing orthodox differentiability.

Definition 5.1 ($\mathcal{R}'_\partial$-linear transformer approximation, **[τ-Effective]**)**.** *Let $f \in \mathsf{Hol}_\tau(X,\mathbb{D})$ and $k \in \mathbb{N}$. We say f is τ-*differentiable at depth k *if the finite-stage stability predicate* Stable_k *certifies that f admits an $\mathcal{R}'_\partial$-linear transformer approximation of bi-degree ≤ 1 in the coordinates (t,x) — i.e. there exists $A_f \in \mathbb{D}$ such that, inside the depth-k witness, increments of f are A_f-linear in the increments of $z = t + jx$.*

Definition **??** is discrete and combinatorial: it is the shadow of differentiation recovered from the registry witness structure, not an imported analytic limit. The algebraic constraint that falls out of "$\mathcal{R}'_\partial$-linearity in z" is the content of Theorem **??** below.

Definition 5.2 (Split Dolbeault operators, **[τ-Effective]**)**.** *Acting on $f = u + jv$ we define the split-complex differential operators*

$$\partial_j \;:=\; \tfrac{1}{2}(\partial_t + j\,\partial_x), \qquad \bar{\partial} \;:=\; \tfrac{1}{2}(\partial_t - j\,\partial_x). \tag{2}$$

Both are $\mathcal{R}'_\partial$-linear operators on functions $X \to \mathbb{D}$.

Note the sign discipline: the split-complex conjugate is $\bar{z} = t - jx$, so $\bar{\partial}$ is the operator $\frac{1}{2}(\partial_t - j\partial_x)$. This parallels the classical complex setup $\partial = \frac{1}{2}(\partial_x - i\partial_y)$, $\bar{\partial} = \frac{1}{2}(\partial_x + i\partial_y)$, but with the algebraic generator j (squaring to $+1$) replacing i (squaring to -1).

5.3 The split-complex Cauchy–Riemann equations

Theorem 5.3 (Split-complex Cauchy–Riemann, **[τ-Effective]**)**.** *Let $f = u + j\,v\colon X \to \mathbb{D}$ with $u, v\colon X \to \mathcal{R}'_\partial$. Then the following are equivalent:*

(i) *f is τ-differentiable at depth k in the sense of Definition* **??** *(equivalently, $f \in \mathsf{Hol}_\tau(X,\mathbb{D})$ at stage k).*

(ii) $\bar{\partial} f = 0$.

(iii) *The real components satisfy*

$$\partial_t u \;=\; \partial_x v, \qquad \partial_x u \;=\; \partial_t v. \tag{3}$$

Proof. The equivalence (i) ⇔ (ii) is the content of extracting the $\mathcal{R}'_\partial$-linear transformer approximation of Definition **??**: an increment $A_f \cdot \Delta z$ with $A_f \in \mathbb{D}$ is exactly the condition that $\bar{\partial} f$ vanishes, since $\bar{\partial} z = 0$ while $\bar{\partial}\bar{z} = 1$.

For (ii) ⇔ (iii), expand using Definition **??**:

$$\bar{\partial} f \;=\; \tfrac{1}{2}(\partial_t - j\partial_x)(u + jv) \;=\; \tfrac{1}{2}(\partial_t u + j\,\partial_t v - j\,\partial_x u - j^2\,\partial_x v).$$

Using $j^2 = +1$ this collapses to

$$\bar{\partial} f \;=\; \tfrac{1}{2}(\partial_t u - \partial_x v) \;+\; \tfrac{j}{2}(\partial_t v - \partial_x u).$$

Because $\{1, j\}$ is an $\mathcal{R}'_\partial$-basis of $\mathbb{D}$, $\bar{\partial} f = 0$ iff both coefficients vanish, giving exactly (**??**). ☐

Remark 5.4 (The load-bearing sign, **[Established]**). The sign structure of (**??**) is the sole load-bearing difference from the classical complex Cauchy–Riemann equations. For $h = u + i\,v$ with $i^2 = -1$, the same computation gives

$$\bar{\partial} h \;=\; \tfrac{1}{2}(\partial_t u - \partial_x v) + \tfrac{i}{2}(\partial_t v + \partial_x u),$$

so vanishing forces $\partial_t u = \partial_x v$ and $\partial_x u = -\partial_t v$. The minus sign in the second classical CR equation comes directly from $i^2 = -1$; in the split-complex case $j^2 = +1$ turns it into a plus sign, and this is exactly the algebraic fingerprint of the wave equation below.

5.4 The wave equation (Main Theorem)

Theorem 5.5 (Wave-Equation Cauchy–Riemann, **[τ-Effective]**). *Let $f \in \mathrm{Hol}_\tau(X, \mathbb{D})$, $f = u + j\,v$. Then each component $u, v\colon X \to \mathcal{R}'_\partial$ satisfies the hyperbolic wave equation*

$$\partial_t^2 u \;=\; \partial_x^2 u, \qquad \partial_t^2 v \;=\; \partial_x^2 v. \tag{4}$$

Proof. By Theorem **??** we have the split-complex CR system (**??**). Apply ∂_t to the first equation $\partial_t u = \partial_x v$:

$$\partial_t^2 u \;=\; \partial_t(\partial_x v) \;=\; \partial_x(\partial_t v),$$

where the last step uses the commutator identity $\partial_t\partial_x = \partial_x\partial_t$. Now substitute the second CR equation $\partial_t v = \partial_x u$:

$$\partial_x(\partial_t v) \;=\; \partial_x(\partial_x u) \;=\; \partial_x^2 u.$$

Combining, $\partial_t^2 u = \partial_x^2 u$. The argument for v is symmetric: apply ∂_t to $\partial_t v = \partial_x u$, then substitute $\partial_t u = \partial_x v$ to obtain $\partial_t^2 v = \partial_x^2 v$. ☐

Remark 5.6 (Elliptic contrast, **[Established]**). The same derivation with $i^2 = -1$ produces the opposite sign. From the classical CR pair $\partial_t u = \partial_x v$, $\partial_x u = -\partial_t v$,

$$\partial_t^2 u \;=\; \partial_x\partial_t v \;=\; \partial_x(-\partial_x u) \;=\; -\partial_x^2 u,$$

which is the Laplace equation $\partial_t^2 u + \partial_x^2 u = 0$. The sign flip is traceable, step by step, to a single swap of $j^2 = +1$ with $i^2 = -1$.

Side-by-side classification of PDE type.. We record the structural dichotomy [**?**] as:

Algebra	Relation	CR equations	Components satisfy
Split-complex $\mathbb{D}$	$j^2 = +1$	$\partial_t u = \partial_x v,\;\; \partial_x u = \partial_t v$	wave: $\partial_t^2 u = \partial_x^2 u$
Elliptic $\mathbb{C}$	$i^2 = -1$	$\partial_t u = \partial_x v,\;\; \partial_x u = -\partial_t v$	Laplace: $\partial_t^2 u + \partial_x^2 u = 0$

The *algebra* determines the *signature of the PDE*: $j^2 = +1$ produces a hyperbolic operator $\partial_t^2 - \partial_x^2$, while $i^2 = -1$ produces an elliptic one $\partial_t^2 + \partial_x^2$. Holomorphy in a boundary algebra is therefore not one thing but two things — and Hinge 4's forcing of $\mathbb{D}$ selects the hyperbolic branch.

The classical several-complex-variables (SCV) theory developed under $i^2 = -1$ is the reference case for the elliptic side; see [**?**, **?**, **?**, **?**] for the modern treatment of holomorphic functions and integral representations in the elliptic setting. The present paper can be read as the hyperbolic $j^2 = +1$ counterpart of that program, restricted from the outset to the boundary algebra $\mathbb{D}$ earned in §**??**.

5.5 Characteristic decomposition

The hyperbolic signature invites a change of coordinates to the null-cone directions, which diagonalises the operator $\partial_t^2 - \partial_x^2$.

Definition 5.7 (Lightcone coordinates, **[τ-Effective]**). *Set*

$$\xi \;:=\; t + x, \qquad \zeta \;:=\; t - x,$$

so that $\partial_\xi = \frac{1}{2}(\partial_t + \partial_x)$ and $\partial_\zeta = \frac{1}{2}(\partial_t - \partial_x)$. Equivalently, in the idempotent basis $z = t + jx = \xi\, e_+ + \zeta\, e_-$.

Proposition 5.8 (e_+/e_--projection depends on a single lightcone coordinate, **[τ-Effective]**). *Let $f = u + jv \in \mathrm{Hol}_\tau(X, \mathbb{D})$, and define the characteristic projections*

$$f_+ \;:=\; u + v, \qquad f_- \;:=\; u - v,$$

so that $f = f_+\, e_+ + f_-\, e_-$. Then

$$\partial_\zeta f_+ \;=\; 0, \qquad \partial_\xi f_- \;=\; 0.$$

Equivalently, f_+ depends only on ξ and f_- depends only on ζ.

Proof. Using (**??**),

$$\partial_\zeta f_+ \;=\; \tfrac{1}{2}(\partial_t - \partial_x)(u + v) \;=\; \tfrac{1}{2}(\partial_t u - \partial_x u + \partial_t v - \partial_x v) \;=\; \tfrac{1}{2}\big((\partial_t u - \partial_x v) + (\partial_t v - \partial_x u)\big) \;=\; 0,$$

each bracket vanishing by one of the two CR equations. The computation for $\partial_\xi f_-$ is symmetric. □

Corollary 5.9 (d'Alembert form of $\mathbb{D}$-holomorphic functions, **[τ-Effective]**). *Every $f \in \mathrm{Hol}_\tau(X, \mathbb{D})$ decomposes uniquely as*

$$f(t, x) \;=\; F(\xi)\, e_+ \;+\; G(\zeta)\, e_- \quad \textit{for some tail-parameterised } F, G, \tag{5}$$

or equivalently, after re-collecting in the $\{1, j\}$ basis,

$$f(t, x) \;=\; \tfrac{1}{2}\big(F(\xi) + G(\zeta)\big) \;+\; j \cdot \tfrac{1}{2}\big(F(\xi) - G(\zeta)\big).$$

Here F, G are tail-parameterised *sector functions, not continuous interior functions in the orthodox sense: they are the $\sim$-stable single-variable transformers arising from Proposition* **??***, evaluated on the light-cone coordinates $\xi, \zeta \in \mathcal{R}'_\partial$ through the boundary ring's $\mathcal{R}'_\partial$-module structure. In particular, Corollary* **??** *is a* boundary-algebra *statement: the interior-point construction that would make F, G classical functions of real variables is deferred to Book II.*

Proof. Combine Proposition **??** with the idempotent decomposition $f = f_+\, e_+ + f_-\, e_-$; set $F := f_+$ and $G := f_-$. □

Remark 5.10 (Hyperbolic vs. complex holomorphy, **[Established]**). In the classical complex case, a holomorphic function is rigidly determined by a *single* function of one complex variable, $h = h(z)$. In the split-complex case, (**??**) shows that a $\mathbb{D}$-holomorphic function is determined by *two independent* functions, one on each light cone. This is the characteristic decomposition in the sense of classical PDE theory, and it is precisely the general solution of the wave equation $\partial_t^2 u = \partial_x^2 u$ in d'Alembert form: $u(t, x) = \frac{1}{2}(F(t + x) + G(t - x))$. The two independent lightcone branches are exactly the e_+ and e_- channels of the boundary algebra.

5.6 Why $j^2 = +1$ gives wave, and why diagonal discipline protects it

Two conceptual takeaways organise the preceding computation.

(a) Hyperbolic signature is a first-principles consequence of $j^2 = +1$.. Theorem **??** does not invoke any choice external to the algebra: the wave equation is forced by the single structural relation $j^2 = +1$ inherited from Hinge 4. There is no parameter to tune, no metric to impose; the signature of the PDE is already present in the multiplication table of the boundary transformer. The Minkowski norm $z\bar{z} = t^2 - x^2$ of (**??**) and the d'Alembert decomposition (**??**) are two faces of the same algebraic fact.

(b) Diagonal discipline protects the hyperbolic branch.. The diagonal discipline developed in §?? forbids re-encoding derivatives as slices through a Cartesian product $X \times \mathbb{D}$ via function graphs. Any such re-encoding would covertly restore the standard complex base and collapse the algebra toward $i^2 = -1$. Under that collapse, the same formal derivation would run with an opposite sign (Remark ??) and yield the Laplace equation instead of the wave equation. The wave-equation Cauchy–Riemann theorem is therefore *protected* by exactly the discipline that forces the codomain $\mathbb{D}$ in the first place; the two commitments — earned codomain (§??) and no Cartesian-product smuggling (§??) — are not independent riders on the theorem but the conditions that secure its conclusion.

5.7 Registry note and Navier–Stokes forward outlook

Registry entries.. Theorem ?? is recorded under II.T (split-CR-wave) in the Book II registry, with the characteristic decomposition (Corollary ??) as II.P (dalembert) [?]. The PDE-type classification table (split vs elliptic) is cross-referenced in Book II Ch. 61.

Forward outlook: τ-Navier–Stokes.. The wave-equation Cauchy–Riemann theorem is the algebraic PDE type underlying the τ-framework treatment of the Navier–Stokes regularity problem in Book III [?]. Because holomorphy in the earned codomain $\mathbb{D}$ enforces a hyperbolic signature, NS dynamics are natively *conservative* at the transformer level: the evolution is a local Hartogs continuation into the light-cone characteristic structure (Corollary ??) rather than a parabolic relaxation toward equilibrium. The regularity question, in that setting, becomes a stabilisation theorem for the e_+/e_- channels (cf. §?? on idempotent-supported sections) rather than a continuation hypothesis requiring external energy estimates. This reframing, which Book III develops in detail, is possible only because the hyperbolic signature is already forced at the holomorphic level, by $j^2 = +1$.

6. DIAGONAL DISCIPLINE: THE NO-CARTESIAN-PRODUCT PRINCIPLE

6.1 The diagonal in classical mathematics: what it does

Before we can state what the *diagonal discipline* forbids, we must remind ourselves how deeply free diagonals penetrate the classical foundations. They are so ubiquitous — and so cheap — that they feel invisible, and this very invisibility is why removing them looks at first like amputation rather than surgery.

Four sites show the phenomenon.

(a) The diagonal morphism in set theory. Given a set X, the map

$$\Delta\colon X \longrightarrow X \times X, \qquad \Delta(x) = (x, x),$$

is free: it exists by the universal property of the Cartesian product, applied to $\mathrm{id}_X, \mathrm{id}_X\colon X \to X$. No axiom is invoked beyond pairing and comprehension. Copying a token of X costs nothing.

(b) Function graphs as slices. A classical function $f\colon X \to Y$ is defined as a subset

$$\Gamma_f \;=\; \{\, (x, y) \in X \times Y \,|\, y = f(x) \,\} \;\subseteq\; X \times Y,$$

i.e. as a *slice through the Cartesian product* $X \times Y$. This definition requires (i) the product object $X \times Y$ to exist as a carrier in its own right, (ii) comprehension to cut out the graph, and (iii) equality of coordinates to be meaningful as a logical predicate. All three are free in ZFC.

(c) Cantor–Gödel diagonalisation. The most famous structural uses of the diagonal — Cantor's proof that $|X| < |\mathcal{P}(X)|$, Gödel's incompleteness, Turing's halting argument — all pivot on passing from a two-variable relation $R(x, y)$ to the one-variable relation $R(x, x)$, i.e. on pulling back along $\Delta\colon X \to X \times X$. The argument is genuinely structural: it is the diagonal itself, not any particular encoding, that generates the paradox.

(d) Contraction in sequent calculus. In classical and intuitionistic sequent calculi, the structural rule

$$\frac{\Gamma, A \vdash B}{\Gamma, A, A \vdash B} \qquad \text{(contraction)}$$

allows one to treat two occurrences of a hypothesis as one, i.e. to *copy* a proof-token. Contraction is precisely the proof-theoretic avatar of the diagonal Δ: in the Curry–Howard correspondence, contraction is interpreted by the diagonal of the ambient Cartesian category.

Remark 6.1 (The underlying operation). What unifies (a)–(d) is a single structural operation: *free copying of a typed token*. In set theory it is Δ, in logic it is contraction, in category theory it is the fact that every Cartesian category has a comonoidal structure $X \to X \otimes X$ for $\otimes = \times$. These are four faces of one move, and classical mathematics grants it for free.

The τ-kernel does not. This is *not* a local stylistic choice but a structural commitment inherited from Book I [?], and its consequences reach all the way to the split-complex signature $j^2 = +1$ of the earned codomain $\mathbb{D}$ (§??). The purpose of the present section is to make that reach fully explicit.

6.2 The τ-kernel discipline: no free diagonal, no free Cartesian product

We now state the discipline. The prohibitions are formulated at the level of admissible transformers and their carriers; they are not commentary, they are structural axioms of the framework.

Definition 6.2 (The Diagonal Discipline). *An admissible τ-framework in the sense of §?? is required to satisfy the following four prohibitions:*

(DD1) No primitive Cartesian product of carriers. *Given admissible carriers X, Y in the framework, the framework does* not *automatically include $X \times Y$ as an admissible carrier. A carrier on which two independent coordinates can be simultaneously and freely projected must be* earned, *not postulated.*

(DD2) No function-graph-as-slice. *An admissible transformer f is not modelled as a subset $\Gamma_f \subseteq X \times Y$ cut out by a graph condition. Transformers are ω-germ data (§??), not graph-predicates on a product.*

(DD3) No exponential object as carrier. *The exponential Y^X of transformers $X \to Y$ is not an admissible carrier. Classical adjunction* $\mathrm{Hom}(X \times Y, Z) \cong \mathrm{Hom}(X, Z^Y)$ *is unavailable because its left-hand side is already unavailable by (DD1).*

(DD4) No free contraction in the meta-logic. *In any meta-logical ledger of a τ-statement, each typed token appears with declared multiplicity; a token cannot be copied silently. The inference $\Gamma, A, A \vdash B \Rightarrow \Gamma, A \vdash B$ requires an explicit structural justification.*

Remark 6.3 (What replaces pairs). (DD1)–(DD4) do not say *"there are no pairs."* They say pairs cannot be had by applying the comprehension schema to a Cartesian product. Inside τ, "pairs" are *typed orbit slots* carried by the four canonical generators $\alpha, \pi, \gamma, \eta$ of the addressability structure: what looks like a pair (x, y) is, after unfolding, a colimit datum bound to *two distinct slots of the same orbit*, not a tuple in a Cartesian product. The full development of this typed-orbit replacement for Cartesian products is the subject of Hinge 7 (forthcoming); for the present paper it suffices that (DD1)–(DD4) hold, because the proof of the main theorem below uses only these hypotheses.

Remark 6.4 (Scope of (DD1)–(DD4)). The four clauses are not independent. (DD1) implies (DD2) and (DD3) immediately: if $X \times Y$ is not a carrier, neither a subset of it nor an exponential into Y with product-adjoint can be. (DD4) is strictly stronger than (DD1)–(DD3): one could imagine a framework without Cartesian products that nevertheless permitted free contraction on certain derived objects. The proof of Theorem ?? below uses all four; we list them separately because each plays a structurally different role.

What we emphasise is the *direction of explanation*. These prohibitions are not postulated ad hoc to engineer the split-complex signature. They flow from the constitutive decision that τ-holomorphy is *prior* to transformers: if transformers are primitive (ω-germ data), there is no meta-stage at which one can first take products and then *cut out* a transformer as a subset of the product. The product-and-slice construction presupposes a classical set-theoretic stratum that the τ-kernel does not recognise.

6.3 The diagonal discipline theorem

We can now state and prove the main structural result of this section.

Theorem 6.5 (Diagonal Discipline, **[τ-Effective]**). *Within the ω-germ transformer framework of §§??–??, there is no admissible construction that produces a function-graph-as-slice $\Gamma_f \subseteq X \times Y$ for a transformer f. Consequently, the idempotent collapse that would force the scalar codomain toward elliptic $\mathcal{R}'_\partial[i]/(i^2+1)$ is structurally precluded, and the split-complex signature $j^2 = +1$ of the earned codomain $\mathbb{D}$ (§??) is preserved. In short: the diagonal-free discipline (DD1)–(DD4) is the* exact *structural obstruction that protects the signature $j^2 = +1$.*

Proof. We argue by contradiction. Suppose for sake of argument that the framework admitted, for every admissible transformer $f\colon X \to Y$, a graph-as-slice representation $\Gamma_f \subseteq X \times Y$.

Step 1 (carrier production). The very statement "$\Gamma_f \subseteq X \times Y$" requires $X \times Y$ to be an admissible carrier of the framework: a subset is a subset of something. This directly violates (DD1).

Step 2 (diagonal is free). Once $X \times Y$ is an admissible carrier for every pair X, Y, the diagonal $\Delta\colon X \to X \times X$ is admissible: it is the unique morphism into $X \times X$ whose two projections are both id_X, and admissibility is closed under universal-property constructions on admissible carriers. Hence *free copying of tokens of* X is available, violating (DD4) via the Curry–Howard translation of Remark **??**.

Step 3 (diagonal resonance). Combine three ingredients now on the table: free contraction (Step 2), equality-as-congruence in the meta-logic (a standard meta-move, not by itself blocked), and the self-product carrier $X \times X$ (Step 1). These are the three ingredients of what we call *diagonal resonance*: the situation where meta-level substitutivity — "x equals x, so everything true of the first x is true of the second" — becomes a *productive* rule at the object level, because the two x's sit in distinct coordinates of a product and can be separately interrogated.

Step 4 (integral-domain force). Diagonal resonance forces the scalar codomain S of the framework to be an integral domain. The argument is constraint-propagation. Suppose S had a zero divisor $a \cdot b = 0$ with $a, b \neq 0$. Consider the transformers $g\colon X \to S$ with $g \equiv a$ and $h\colon X \to S$ with $h \equiv b$ (constants valued in S). Their product transformer $g \cdot h$ has graph $\Gamma_{g \cdot h} = \{(x, 0) : x \in X\}$. But by the graph-of-product universal property (Step 1 makes this available), $\Gamma_{g \cdot h}$ must agree with the image of $\Gamma_g \times_X \Gamma_h$ under the multiplication-on-S map — and this image, unfolded via Step 2's free diagonal, discriminates between the two input copies of x, so the resulting slice is not well-defined as a subset of $X \times S$. The only way to repair coherence is to outlaw zero divisors in S, i.e. to require S to be an integral domain.

Step 5 (elliptic forcing). By Theorem 1.7 of [**?**] (Hinge 4, the Elliptic Complex Exclusion), if S is a two-dimensional commutative algebra over $\mathbb{R}$ of the form $\mathcal{R}'_\partial[i]/(i^2 - c)$ with $c \in \{-1, 0, +1\}$, then: the parabolic case $c = 0$ yields $\mathcal{R}'_\partial[\epsilon]/(\epsilon^2)$, which has nilpotents and is not an integral domain; the hyperbolic case $c = +1$ yields $\mathcal{R}'_\partial[j]/(j^2 - 1) = \mathbb{D}$, which has zero divisors $e_+ = (1 + j)/2$ and $e_- = (1 - j)/2$ and is not an integral domain. Only the elliptic case $c = -1$, giving $\mathcal{R}'_\partial[i]/(i^2 + 1) \cong \mathbb{C}$, is an integral domain. So Step 4 forces $S \cong \mathbb{C}$ in the two-dimensional setting.

Step 6 (idempotent annihilation). But Theorem 1.7 of [**?**] also records that $\mathbb{C}$ has no nontrivial idempotents: the equation $z^2 = z$ has only the solutions $z = 0, 1$. The swapped idempotent pair (e_+, e_-) that carries the B/C sector decomposition of Hinge 3 [**?**] *does not survive* in $\mathbb{C}$. Hence the σ-equivariant split required by the boundary algebra ceases to exist, and the framework collapses.

Step 7 (contradiction). The framework of §§**??**–**??** manifestly exhibits the σ-equivariant idempotent pair (e_+, e_-) — it is the very structure established in §**??**. Hence the hypothesis of Step 1 is false: *no* admissible construction within the framework produces graphs-as-slices. Equivalently, (DD1)–(DD4) hold, and the split-complex signature $j^2 = +1$ is preserved. □

Remark 6.6 (On the logical status of Theorem **??**). The theorem is tagged **[τ-Effective]** because, while the algebraic core of the proof is pure commutative algebra plus meta-logic, the load-bearing Step 4 (the forcing of integral-domain structure via diagonal resonance) invokes the τ-kernel's coherence predicate as an irreducible input, with a full Step-4 lemma promised in the forthcoming Hinge 7 canonical-addressability paper. The theorem also depends on the sector-decomposition fact of §**??**. In particular, the theorem would hold verbatim for any framework in which (i) (DD1)–(DD4) are imposed and (ii) a σ-equivariant idempotent split exists. The *surprise* content is Step 4 — that graphs-as-slices *force* integral-domain scalars — which we isolate as a stand-alone lemma in Hinge 7.

6.4 Three faces of the same move

The three radical inversions that structure this paper are not independent moves; they are the same structural decision viewed from three angles. We record this explicitly.

Theorem 6.7 (Equivalence of the Three Inversions, **[τ-Effective]**). *For an admissible τ-scalar-codomain S of the framework of §§**??**–**??**, the following are mutually equivalent:*

(I1) *(DD1)–(DD4) hold in the ambient framework; i.e.* no function-graph-as-slice through a Cartesian product is admissible.

(I2) *S carries the split-complex signature $j^2 = +1$; equivalently, $S \cong \mathbb{D}$ over $\mathcal{R}'_\partial$.*

(I3) *The Cauchy–Riemann equations for admissible transformers $X \to S$ decouple into the hyperbolic wave equation of §??, rather than Laplace's equation.*

Proof sketch. The implications (I1)⇒(I2)⇒(I3) have already been established: (I1)⇒(I2) is Theorem **??**, and (I2)⇒(I3) is Theorem 5.4 (the wave-equation Cauchy–Riemann, §**??**), where the sign of j^2 enters as the sign between ∂_x^2 and ∂_y^2 in the decoupled second-order equation.

For (I3)⇒(I1): if the admissible second-order operator on S-valued transformers is the hyperbolic wave operator $\partial_x^2 - \partial_y^2$, then the characteristic cones of the operator give two independent null directions $x \pm y$, which force a two-dimensional *split* of S into null-eigenbundles $e_+ \cdot S \oplus e_- \cdot S$. This split is only compatible with scalar multiplication if S contains nontrivial idempotents, i.e. if S is not an integral domain. By the contrapositive of Step 4 of the proof of Theorem **??**, S not an integral domain requires the failure of graph-as-slice admissibility, which is (DD1)–(DD4). The detailed sign-tracking is deferred to Hinge 7, where it becomes one case of the broader Hinge 7 classification. □

Remark 6.8 (What the equivalence buys us). Theorem **??** says that the author of a τ-holomorphy paper has exactly one choice to make, not three. Once any one of (I1), (I2), (I3) is fixed, the other two follow. This is why the paper can speak of *"the"* diagonal discipline, *"the"* split-complex signature, *"the"* wave-Cauchy–Riemann: up to reformulation, they are one object.

6.5 Structural moral: why this is the right discipline

The pedagogical temptation, on first encountering (DD1)–(DD4), is to read them as *restrictions* — as things the τ-framework is forbidden to do relative to the classical baseline — and therefore to ask what compensating axiom is being added elsewhere to make up the loss. This reading is backward.

A direct parallel is *linear logic*. Girard's linear logic forbids free contraction and weakening in its structural rules. At first glance this looks crippling: one cannot copy hypotheses, one cannot discard them. But what emerges is a logic with strictly richer resource-tracking capacity than classical or intuitionistic logic — a logic in which quantum-mechanical entanglement, concurrent processes, and proof nets find their natural home. The no-free-contraction restriction *creates* a discipline that buys genuine new expressive power, and classical logic is recovered as the contraction-and-weakening quotient sitting above.

The diagonal discipline (DD1)–(DD4) plays the same role for τ-holomorphy. What it "gives up" — unrestricted Cartesian products, function graphs, free diagonals — is the freedom to collapse structurally distinct objects together through equational reasoning. What it "gains" is the ability to preserve the rich asymmetry of Hinge 3's σ-involution and Hinge 4's boundary algebra all the way into the codomain: the swapped idempotent pair (e_+, e_-) of $\mathbb{D}$ survives precisely because there is no meta-logical pressure collapsing it into the two constant classes $\{0, 1\}$ of a would-be elliptic $\mathbb{C}$.

This is the point of Hinge 4's own emphasis in [?]: "in split-complex numbers there are exactly two nontrivial idempotents; elliptic complex numbers have no nontrivial idempotents, making the B/C distinction ontically impossible there." Pulled back through Theorem **??**, this reads as: *the B/C distinction is ontically possible in τ-holomorphy precisely because, and exactly to the extent that, the diagonal discipline holds.* The split-complex signature is not a convenient technical choice; it is where the framework's own structural commitments inevitably land.

There is a further, more local observation. Book IV's virtualisation apparatus [?] repeatedly exploits that certain sector decompositions at the bulk level descend to σ-equivariant idempotent splits at the boundary. The present section explains why these descents are not miracles: they descend *because* the ambient framework refuses the diagonal. If graphs-as-slices were admissible, the bulk sector structure would be flattened through the same integral-domain collapse traced in the proof of Theorem **??**, and the descent results of Book IV would fail coherently rather than work coherently.

6.6 Registry and Lean preview

The diagonal-discipline material of this section is tracked in the τ-registry as follows. The four prohibitions (DD1)–(DD4) are registered as axioms under the Book II holomorphy chapter devoted to admissibility; Theorem **??** is registered at scope **[τ-Effective]** (with a forthcoming Hinge-7 Step-4 lemma to sharpen it toward **[Established]**), with dependencies on Hinge 4's elliptic exclusion (Theorem 1.7 of [?]); the three-inversions equivalence (Theorem **??**) is registered at scope **[τ-Effective]**, with the Hinge 7 reference as an open forward dependency.

The Lean formalisation is slated for the module

`TauLib.BookII.Holomorphy.DiagonalDiscipline`, carrying the following definitions and results: the structure `AdmissibleFramework` gathering (DD1)–(DD4); the lemma `graph_slice_forces_integral_domain` implementing Step 4 of the proof; the theorem `diagonal_discipline` implementing Theorem **??**; and the theorem `three_inversions_equivalent` implementing Theorem **??**. The module sits under the earned-codomain module `TauLib.BookII.Holomorphy.EarnedCodomain` (which fixes $\mathbb{D}$) and is imported by the forthcoming earned-category module `TauLib.BookII.Holomorphy.EarnedCategory` that will be the subject of §**??**.

For the τ-topos construction of Hinge 6, the content of Theorem **??** is the hypothesis that lets us *not* be a Cartesian category without thereby being structurally impoverished: the earned monoidal product that replaces the Cartesian one will be constructed in Hinge 6 and shown, via Theorem **??** above, to carry exactly the split-complex signature demanded by the codomain $\mathbb{D}$. In this sense the present section is the pivot of Hinges 4–6: it is where the scalar-level exclusion of Hinge 4 is revealed as a carrier-level discipline, which is then promoted to a topos-level construction in Hinge 6.

7. THE EARNED CATEGORICAL MACHINE

7.1 The setup: what we have and what we don't

The preceding four sections have built, brick by brick, a pre-categorical ontology: ω-tails (§**??**), the certified-transformer type $\mathsf{Hol}_\tau(X, Y)$ (§**??**), the forced scalar codomain $\mathbb{D}$ (§**??**), and the hyperbolic wave-equation Cauchy–Riemann signature (§**??**). Nowhere in that development was any categorical notion *imposed*: we did not postulate a composition operator, an identity morphism, associativity, or any functorial assignment. All we admitted were boundary germs coded as normal-form programs, together with three admissibility predicates (Typed, Stable, tail-independence) that operate on those programs. The present section proves that the full categorical apparatus emerges from exactly this minimal stock, as theorems rather than axioms. This is the sense in which the resulting category is *earned* [?, ?].

We make the "before / after" inventory explicit, because the argument's substance lies in what remains prohibited throughout.

What we have (the primitive stock)..

- **ω-tails** Ω_{tail} (Def. **??**) and the decidable prefix predicates $\mathsf{Pref}_{k,\sigma}$.
- **Tail-equivalence** $\sim$ (Def. **??**), which is an equivalence relation (Lemma **??**).
- **NF-coded transformer codes** $c \in \mathsf{Code}$ with semantics $[\![c]\!] : \Omega_{\text{tail}} \to \Omega_{\text{tail}}$ (Def. **??**).
- **Admissibility predicates** $\mathsf{Typed}(X, Y, c)$, $\mathsf{Stable}(X, Y, c)$, and tail-independence beyond a finite depth (Defs. **??**–**??**).
- **The transformer type** $\mathsf{Hol}_\tau(X, Y)$ as the intersection of those three predicates (Def. **??**).

What we do *not* import..

- No **"category"** axiomatically imposed: neither an abstract composition $\circ$ nor the associativity law nor unit laws are assumed. They will each be proved below.
- No **categorical product** $X \times Y$ or exponential Y^X. The diagonal discipline of §**??** keeps these unavailable, and we shall not sneak them back in.
- No **imposed functor**: the assignment $Y \mapsto \mathsf{Hol}_\tau(X, Y)$ exists set-theoretically in the sense of countable-type-forming operations from Code, but its functoriality must be earned.
- No **ambient morphism language**. We do not assume a notion of "morphism of carriers" beyond what the refinement discipline of tails forces us to admit (§**??**).

Remark 7.1 (Why this bookkeeping matters)**.** Classical constructions of a category of holomorphic maps presuppose that composition exists: one *defines* $(g \circ f)(x) := g(f(x))$ and then verifies that the predicate "holomorphic" is preserved. In our setting there is no pointwise composition available, because there is no pointwise evaluation available: the tails are not functions from points to values but intensional coherence certificates. The question to answer is therefore not "is Hol_τ closed under composition?" but rather "does a binary operation on Hol_τ-representatives exist that deserves the name composition?" Section §**??** shows that the answer is yes and that the operation is canonical (code concatenation modulo NF reduction). The rest of this section then unpacks the consequences.

7.2 Earned composition

Theorem 7.2 (Earned composition, **[τ-Effective]**). *V.T.EC.COMP Let $X, Y, Z \in \mathrm{Obj}(\tau)$. For every $f \in \mathsf{Hol}_\tau(X, Y)$ and $g \in \mathsf{Hol}_\tau(Y, Z)$, the code*

$$c_{g\circ f} \;:=\; \mathsf{Norm}(c_g \cdot c_f),$$

obtained by concatenating the representatives c_f, c_g of f, g and reducing to normal form, satisfies

$$\mathsf{Typed}(X, Z, c_{g\circ f}) \;\wedge\; \mathsf{Stable}(X, Z, c_{g\circ f}) \;\wedge\; \textit{tail-independent beyond some depth}.$$

In particular, $c_{g\circ f}$ represents an element $g \circ f \in \mathsf{Hol}_\tau(X, Z)$. The induced binary operation $\circ\colon \mathsf{Hol}_\tau(Y, Z) \times \mathsf{Hol}_\tau(X, Y) \to \mathsf{Hol}_\tau(X, Z)$ is independent of the choice of $\sim$-representatives.

Proof. Let $c_f, c_g \in \mathsf{Code}$ be chosen NF representatives of f and g. We verify the three admissibility conjuncts on the raw concatenation $c_g \cdot c_f$; normalisation to $c_{g\circ f} = \mathsf{Norm}(c_g \cdot c_f)$ preserves them because Norm is $\sim$-equivalent to the identity on semantics.

Step 1: Typing. Let $t \in \Omega_{\text{tail}}$ with $\mathsf{Tail}_X(t)$. Because $\mathsf{Typed}(X, Y, c_f)$, we have $\mathsf{Tail}_Y([\![c_f]\!](t))$. Because $\mathsf{Typed}(Y, Z, c_g)$, we have

$$\mathsf{Tail}_Z([\![c_g]\!]([\![c_f]\!](t))) \;=\; \mathsf{Tail}_Z([\![c_g \cdot c_f]\!](t)),$$

using the standard semantic identity $[\![c_g \cdot c_f]\!] = [\![c_g]\!] \circ [\![c_f]\!]$ for NF programs. Hence $\mathsf{Typed}(X, Z, c_g \cdot c_f)$ holds, which proves $\mathsf{Typed}(X, Z, c_{g\circ f})$ after normalisation.

Step 2: Stability. Let $t, t' \in \Omega_{\text{tail}}$ with $\mathsf{Tail}_X(t), \mathsf{Tail}_X(t')$ and $t \sim t'$. By $\mathsf{Stable}(X, Y, c_f)$ we have $[\![c_f]\!](t) \sim [\![c_f]\!](t')$. Typing gives Tail_Y for both sides. Applying $\mathsf{Stable}(Y, Z, c_g)$ then yields

$$[\![c_g \cdot c_f]\!](t) \;=\; [\![c_g]\!]([\![c_f]\!](t)) \;\sim\; [\![c_g]\!]([\![c_f]\!](t')) \;=\; [\![c_g \cdot c_f]\!](t'),$$

establishing $\mathsf{Stable}(X, Z, c_g \cdot c_f)$. The equivalence is preserved by NF reduction because $\sim$ is a congruence for Norm.

Step 3: Tail-independence. Suppose c_f is tail-independent beyond depth k_f on $X \to Y$ and c_g is tail-independent beyond depth k_g on $Y \to Z$. Because c_f acts by a finite-depth code (finite witness data attached at each elementary rewriting step), its depth-shift is bounded: there exists $\delta_f \in \mathbb{N}$ such that, for every $t, t' \in \Omega_{\text{tail}}$ with $\mathsf{Tail}_X(t), \mathsf{Tail}_X(t')$ and $t \equiv_k t'$ with $k \geq k_f + \delta_f$, the output pair $[\![c_f]\!](t), [\![c_f]\!](t')$ agrees on the first $k - \delta_f$ depths. In particular, choosing $k_0 := \max(k_f, k_g) + \delta_f$, we have: for every $t \equiv_{k_0} t'$ admissible for X, the images satisfy $[\![c_f]\!](t) \equiv_{k_0 - \delta_f} [\![c_f]\!](t')$, so that $[\![c_f]\!](t) \equiv_{k_g} [\![c_f]\!](t')$. Applying tail-independence of c_g beyond depth k_g gives

$$[\![c_g \cdot c_f]\!](t) \;\sim\; [\![c_g \cdot c_f]\!](t').$$

Therefore $c_g \cdot c_f$ is tail-independent beyond depth k_0; Norm at most *decreases* the witness depth (never increases it), so $c_{g\circ f}$ is tail-independent beyond the same finite depth.

Step 4: $\sim$-representative independence. Let $c_f, \tilde{c}_f$ be two representatives of the same element $f \in \mathsf{Hol}_\tau(X, Y)$ (that is, they induce $\sim$-equal semantics on Tail_X), and similarly $c_g, \tilde{c}_g$ for g. Then $[\![c_g \cdot c_f]\!]$ and $[\![\tilde{c}_g \cdot \tilde{c}_f]\!]$ are $\sim$-equal on Tail_X by the functorial composition of $\sim$-equalities (Steps 1 and 2 used both directions of the stability bound). Normal-form reduction sends them to the same canonical NF code, because Norm is a function on $\sim$-semantic-classes. Hence $c_{g\circ f} = \mathsf{Norm}(c_g \cdot c_f)$ is well-defined on the level of Hol_τ-elements. □

Remark 7.3 (What this theorem does *not* say). Theorem **??** does not assert that "holomorphic maps compose," in the sense of a pre-existing composition operator inherited from a larger category. It asserts that the raw string operation "concatenate codes and normalise" — a purely syntactic procedure on finite program data — *turns out* to preserve the three admissibility predicates. The name "composition" is assigned afterwards, once the procedure has been shown to have the right closure properties. This is the precise technical meaning of "earned" here: we prove closure first, and only then give the operation a categorical name.

7.3 Earned identity

Definition 7.4 (The identity NF code). *V.D.EC.ID For each carrier $X \in \mathrm{Obj}(\tau)$, let $c_{\mathrm{id}_X} \in \mathsf{Code}$ denote the* empty NF code*: the unique code whose underlying rewriting program contains no elementary steps, with semantics*

$$[\![c_{\mathrm{id}_X}]\!](t) \;=\; t \qquad \textit{for all } t \in \Omega_{\mathrm{tail}}.$$

Proposition 7.5 (Identity admissibility, [τ-Effective]). *V.P.EC.IDADM For each carrier X, the empty code c_{id_X} satisfies $\mathsf{Typed}(X, X, c_{\mathrm{id}_X})$, $\mathsf{Stable}(X, X, c_{\mathrm{id}_X})$, and tail-independence beyond depth 0. Hence the element*

$$\mathrm{id}_X \;:=\; c_{\mathrm{id}_X} \;\in\; \mathsf{Hol}_\tau(X, X)$$

is well-defined.

Proof. Typed: trivial, since $[\![c_{\mathrm{id}_X}]\!](t) = t$ leaves $\mathsf{Tail}_X(t)$ unchanged. Stable: if $t \sim t'$, then $[\![c_{\mathrm{id}_X}]\!](t) = t \sim t' = [\![c_{\mathrm{id}_X}]\!](t')$. Tail-independence beyond depth 0: for all $k \geq 0$, if $t \equiv_k t'$ then $[\![c_{\mathrm{id}_X}]\!](t) = t \equiv_k t' = [\![c_{\mathrm{id}_X}]\!](t')$, which is certainly $\sim$. □

Proposition 7.6 (Unit laws, [τ-Effective]). *V.P.EC.UNIT For every $f \in \mathsf{Hol}_\tau(X, Y)$,*

$$\mathrm{id}_Y \circ f \;=\; f \qquad \textit{and} \qquad f \circ \mathrm{id}_X \;=\; f$$

in $\mathsf{Hol}_\tau(X, Y)$, where $\circ$ is the composition of Theorem ??.

Proof. Concatenating any code c_f with the empty code on either side yields c_f as a raw string: $c_{\mathrm{id}_Y} \cdot c_f = c_f = c_f \cdot c_{\mathrm{id}_X}$. Normalisation is idempotent on NF codes, so $\mathsf{Norm}(c_{\mathrm{id}_Y} \cdot c_f) = \mathsf{Norm}(c_f) = c_f$ and likewise on the right. Hence the Hol_τ-element represented is f in both cases. □

Remark 7.7 (Representatives vs. equivalence classes). Proposition ?? is stated at the level of NF codes, where the unit laws hold on the nose. Seen through the $\sim$-semantic quotient, they hold in the obvious sense of Hol_τ-equality. No categorical weakening to isomorphism is needed: the identity transformer is strictly a two-sided unit for $\circ$ on representatives.

7.4 Earned associativity

We come to the associativity law, which in classical category theory is posed as an axiom. In our setting it must be proved, and its proof rests on two independent features of the NF-code framework: the string associativity of concatenation, and the Church–Rosser (confluence) property of the τ-native rewriting system on codes.

Theorem 7.8 (Earned associativity, [τ-Effective]). *V.T.EC.ASSOC For every quadruple of carriers $X, Y, Z, W \in \mathrm{Obj}(\tau)$ and every composable triple*

$$f \in \mathsf{Hol}_\tau(X, Y), \qquad g \in \mathsf{Hol}_\tau(Y, Z), \qquad h \in \mathsf{Hol}_\tau(Z, W),$$

the composites $(h \circ g) \circ f$ and $h \circ (g \circ f)$ in $\mathsf{Hol}_\tau(X, W)$ agree:

$$(h \circ g) \circ f \;=\; h \circ (g \circ f).$$

Proof. At the level of raw code strings, concatenation is associative as a binary operation:

$$(c_h \cdot c_g) \cdot c_f \;=\; c_h \cdot (c_g \cdot c_f) \qquad \text{as elements of } \mathsf{Code}.$$

This is a purely syntactic observation about finite sequences of elementary rewriting steps. Applying Norm to both sides gives

$$\mathsf{Norm}((c_h \cdot c_g) \cdot c_f) \;=\; \mathsf{Norm}(c_h \cdot (c_g \cdot c_f)),$$

because Norm is a function of the underlying raw string.

It remains to relate the two composites of Theorem **??** to the raw-string associativity. By definition of $\circ$,

$$(h \circ g) \circ f \;=\; \mathsf{Norm}(\mathsf{Norm}(c_h \cdot c_g) \cdot c_f),$$

$$h \circ (g \circ f) \;=\; \mathsf{Norm}(c_h \cdot \mathsf{Norm}(c_g \cdot c_f)).$$

Confluence of the τ-native rewriting system on Code (Church–Rosser; see [**?**] and the canonical-addressability infrastructure of Hinge 7 for the detailed verification) asserts: for every raw concatenation, repeated normalisation reaches a *unique* canonical form, independent of the order in which intermediate normalisations are performed. Specifically, the two iterated-normalisation expressions above reduce to the same NF code as the single-pass normalisation $\mathsf{Norm}(c_h \cdot c_g \cdot c_f)$:

$$\mathsf{Norm}(\mathsf{Norm}(c_h \cdot c_g) \cdot c_f) \;=\; \mathsf{Norm}(c_h \cdot c_g \cdot c_f) \;=\; \mathsf{Norm}(c_h \cdot \mathsf{Norm}(c_g \cdot c_f)).$$

Hence $(h \circ g) \circ f$ and $h \circ (g \circ f)$ are equal as Hol_τ-elements. □

Remark 7.9 (Two inputs, one output)**.** The proof uses two features independently. *Syntactic* associativity of raw concatenation is trivial. *Semantic* associativity — the fact that iterated normalisation commutes with concatenation — is not trivial; it is precisely the Church–Rosser property of the NF rewriting system. Without confluence the earned category would fail associativity, even though classical composition of functions is trivially associative. This is the price of doing without an ambient set-theoretic function space: the equational theory must be verified on codes, and the equality of compositional outcomes is an *analytic* property of the rewriting dynamics, not a definitional truism. The canonical-addressability development of Hinge 7 supplies this verification; here we cite it as an input and observe that associativity in Hol_τ is its first categorical dividend.

7.5 Earned functoriality

Having established composition, identity, associativity, and the unit laws, we can now speak meaningfully of the category $\mathbf{Hol}_\tau$ (§**??** below). The next question is functoriality: does the assignment $X \mapsto \mathsf{Hol}_\tau(X, -)$ extend to a structure-preserving map out of a suitable category of carriers?

To answer, we must first identify the category from which the functorial input is drawn. We call it the *probe category* P_τ, anticipating the terminology of Hinge 6.

Definition 7.10 (The probe category P_τ)**.** *V.D.EC.PROBE The* probe category P_τ *has:*

- ***Objects:*** *carriers* $X \in \mathrm{Obj}(\tau)$.
- ***Morphisms:*** probe refinements $\rho\colon X \to X'$*, given by NF codes* $c_\rho \in \mathsf{Code}$ *that preserve typing (*$\mathsf{Typed}(X, X', c_\rho)$*), preserve stability (*$\mathsf{Stable}(X, X', c_\rho)$*), are tail-independent beyond finite depth, and additionally satisfy the* boundary-preserving closure *condition: for each admissible* $t \in \Omega_{\mathrm{tail}}$ *with* $\mathsf{Tail}_X(t)$*, the image* $[\![c_\rho]\!](t)$ *is a finer-cylinder representative at each primorial depth (cf. [**?**] for the primorial filtration).*

*Identity morphisms and composition are inherited from §§**??**–**??**.*

Probe morphisms are themselves a special class of ω-germ transformers, distinguished by the boundary-preserving closure condition. They are the τ-native analogue of inclusions of cylinder bases in the classical profinite picture.

Theorem 7.11 (Earned functoriality, **[τ-Effective]**)**.** *V.T.EC.FUNCT Fix* $Y \in \mathrm{Obj}(\tau)$*. The assignment*

$$X \;\longmapsto\; \mathsf{Hol}_\tau(X, Y), \qquad (\rho\colon X \to X') \;\longmapsto\; (\rho^*\colon \mathsf{Hol}_\tau(X', Y) \to \mathsf{Hol}_\tau(X, Y)),$$

with $\rho^*(c_{f'}) := \mathsf{Norm}(c_{f'} \cdot c_\rho)$*, defines a contravariant functor from* P_τ *to the category of countable types. That is:*

(F1) $\mathrm{id}_X^* = \mathrm{id}_{\mathsf{Hol}_\tau(X,Y)}$.

(F2) $(\rho' \circ \rho)^* = \rho^* \circ \rho'^*$ *for every composable pair* $X \xrightarrow{\rho} X' \xrightarrow{\rho'} X''$ *in* P_τ*.*

(F3) *The boundary-preserving closure of* ρ *guarantees that* ρ^* *maps* $\mathsf{Hol}_\tau(X', Y)$ *into* $\mathsf{Hol}_\tau(X, Y)$*: every admissible transformer pulled back along a probe refinement is again admissible.*

Proof. *(F1) Identity preservation.* If $\rho = \mathrm{id}_X$, then c_ρ is the empty NF code (Def. **??**). For any $c_{f'} \in \mathsf{Hol}_\tau(X, Y)$,

$$\rho^*(c_{f'}) \;=\; \mathsf{Norm}(c_{f'} \cdot c_{\mathrm{id}_X}) \;=\; c_{f'},$$

by the right unit law (Prop. **??**). Hence $\mathrm{id}_X^* = \mathrm{id}_{\mathsf{Hol}_\tau(X,Y)}$.

(F2) Contravariance. Given $\rho\colon X \to X'$ and $\rho'\colon X' \to X''$, set $\rho' \circ \rho\colon X \to X''$ by code concatenation $c_{\rho'\circ\rho} = \mathsf{Norm}(c_{\rho'} \cdot c_\rho)$. For any $c_{f''} \in \mathsf{Hol}_\tau(X'', Y)$,

$$\begin{aligned}(\rho' \circ \rho)^*(c_{f''}) &= \mathsf{Norm}(c_{f''} \cdot \mathsf{Norm}(c_{\rho'} \cdot c_\rho)) \\ &= \mathsf{Norm}(c_{f''} \cdot c_{\rho'} \cdot c_\rho) \qquad \text{(confluence)} \\ &= \mathsf{Norm}(\mathsf{Norm}(c_{f''} \cdot c_{\rho'}) \cdot c_\rho) \qquad \text{(confluence)} \\ &= \rho^*(\rho'^*(c_{f''})) = (\rho^* \circ \rho'^*)(c_{f''}).\end{aligned}$$

The middle two equalities invoke the same Church–Rosser property used in Theorem **??**.

(F3) Admissibility preservation. We check that $\rho^*(c_{f'}) \in \mathsf{Hol}_\tau(X, Y)$ when c_ρ is a probe refinement and $c_{f'} \in \mathsf{Hol}_\tau(X', Y)$. Typing: $\mathsf{Typed}(X, X', c_\rho)$ and $\mathsf{Typed}(X', Y, c_{f'})$ together give $\mathsf{Typed}(X, Y, c_{f'} \cdot c_\rho)$, preserved by Norm. Stability: by the same two-step argument as Step 2 of Theorem **??**. Tail-independence: tail-independence of c_ρ beyond depth k_ρ and of $c_{f'}$ beyond depth $k_{f'}$ yields tail-independence of the composite beyond $\max(k_\rho, k_{f'}) + \delta_\rho$, where δ_ρ is the finite depth-shift of ρ. The *boundary-preserving closure* condition on ρ (Def. **??**) is exactly what ensures $\delta_\rho < \infty$: without it, ρ could introduce arbitrary primorial-depth drift and tail-independence bounds could blow up. With it, the composition has a finite tail-independence depth, as required. □

Remark 7.12 (The naturality condition is boundary-preserving closure)**.** Theorem **??** makes the naturality condition for $X \mapsto \mathsf{Hol}_\tau(X, -)$ explicit: it is the *boundary-preserving closure* of the morphism ρ in P_τ. In classical sheaf theory, the analogous role is played by the requirement that restriction maps be sheaf morphisms. In our setting, without ambient continuity, the requirement is intrinsically ω-combinatorial: finite tail-independence must survive pullback. Remarkably, this single condition — a finite depth bound being preserved under refinement — replaces the entire classical analytic infrastructure (continuity, compactness, partition-of-unity arguments) in securing functoriality. The economy is characteristic of the τ-framework [**?**].

Remark 7.13 (Covariant version)**.** The covariant counterpart — $Y \mapsto \mathsf{Hol}_\tau(X, Y)$ with post-composition along morphisms of P_τ on the target — works by the same argument with the factors in the opposite order: if $\sigma\colon Y \to Y'$ is a probe refinement, then $\sigma_*\colon \mathsf{Hol}_\tau(X, Y) \to \mathsf{Hol}_\tau(X, Y')$ is defined by $\sigma_*(c_f) := \mathsf{Norm}(c_\sigma \cdot c_f)$, and the same boundary-preserving closure assumption yields covariant functoriality. Both variants will be used in Hinge 6's τ-topos development (forthcoming).

7.6 What we have built: the category $\mathbf{Hol}_\tau$

The earned composition, identity, associativity, and unit laws combine into a single synthetic statement.

Corollary 7.14 (The earned category $\mathbf{Hol}_\tau$, **[τ-Effective]**)**.** *V.C.EC.CAT The data*

$$\mathbf{Hol}_\tau \;:=\; \big(\mathrm{Obj}(\tau),\; (X, Y) \mapsto \mathsf{Hol}_\tau(X, Y),\; \circ,\; \mathrm{id}_{(-)}\big)$$

forms a category, with:

- ***Objects:*** *carriers $X \in \mathrm{Obj}(\tau)$.*
- ***Morphisms:*** *admissible ω-germ transformers $\mathsf{Hol}_\tau(X, Y)$ (Def. **??**).*
- ***Composition:*** *code concatenation modulo NF reduction (Theorem **??**).*
- ***Identity:*** *the empty NF code (Prop. **??**).*
- ***Associativity & unit laws:*** *earned as theorems (Theorem **??** and Prop. **??**).*

*The inclusion $\mathbf{Hol}_\tau \hookrightarrow \mathbf{Cat}_\tau$ into the full τ-topos [**?**] is a full subcategory inclusion.*

Proof. Direct combination of Theorems **??**, **??**, and **??**, together with Propositions **??** and **??**. The full-subcategory embedding into $\mathbf{Cat}_\tau$ follows from the universal property of certified transformers in the τ-topos, developed in Hinge 6 (forthcoming). □

Remark 7.15 (The inversion)**.** We have constructed a category, but the order of construction is inverted compared to classical practice. Classically, one begins with a category (sets, topological spaces, manifolds, schemes, ...) and then singles out the holomorphic morphisms as a distinguished class. Here we began with the notion of ω-germ transformer — holomorphy itself — and constructed the category around it *a posteriori*. The holomorphic morphisms are not a substructure of a larger category; they are the primitive, and the category is their derived shadow. This is the structural sense in which $\mathbf{Hol}_\tau$ is an *earned* category: no categorical axiom was assumed; every piece of categorical structure was proved as a theorem of the ω-germ admissibility predicates.

7.7 Relation to other τ-categories (forward outlook)

The earned category $\mathbf{Hol}_\tau$ sits inside the larger architecture of τ-categorical objects. We sketch the connections here; precise constructions appear in subsequent work.

$\mathbf{Hol}_\tau \subset \mathbf{Cat}_\tau$.. The full τ-topos $\mathbf{Cat}_\tau$ over the probe category P_τ is developed in Hinge 6 (forthcoming). Its morphisms include all admissible transformers, of which $\mathsf{Hol}_\tau(X, Y)$ is the boundary-first subtype. The inclusion $\mathbf{Hol}_\tau \hookrightarrow \mathbf{Cat}_\tau$ is full and faithful by construction; the key structural fact is that the subobject classifier Ω_τ of $\mathbf{Cat}_\tau$ sees $\mathbf{Hol}_\tau$ as a principal σ-equivariant sub-topos. The technical details of this embedding are beyond the scope of the present paper.

Endomorphism subcategory HolEnd_τ.. Restricting the object arguments to coincide, $X = Y$, gives the monoid-valued endomorphism structure

$$\mathsf{HolEnd}_\tau \;:=\; \{(X, f)\colon X \in \mathsf{Obj}(\tau),\;\; f \in \mathsf{Hol}_\tau(X, X)\}.$$

This is developed in §**??**, where HolEnd_τ is shown to be the first τ-category on which the pre-Yoneda collapse operates, and where the idempotent-supported holomorphy theorem of §**??** reveals its characteristic B/C-bipartite structure. The identity and composition of HolEnd_τ are inherited from $\mathbf{Hol}_\tau$ and hence are themselves earned rather than imposed.

Book II's holomorphic endomorphism category.. Book II [**?**] develops the full theory of holomorphic endomorphism categories (the earned-composition and earned-category programmes, building on Book I's normal-form machinery) and recovers the structural results of this section as a prerequisite to its Central Theorem $\mathcal{O}(\tau^3) \cong A_{\text{spec}}(\mathbb{L})$. The present paper provides the boundary-algebra restriction of that development; the interior-point extension via Hartogs continuation is the next step.

7.8 Registry and Lean preview

Remark 7.16 (Registry identifiers, **[τ-Effective]**)**.** The theorems of this section will be registered under `registry/book2_registry.jsonl`:

- Theorem **??** (earned composition): V.T.EC.COMP.
- Proposition **??** (identity admissibility): V.P.EC.IDADM.
- Proposition **??** (unit laws): V.P.EC.UNIT.
- Theorem **??** (earned associativity): V.T.EC.ASSOC.
- Theorem **??** (earned functoriality): V.T.EC.FUNCT.
- Corollary **??** (the earned category $\mathbf{Hol}_\tau$): V.C.EC.CAT.

These collectively establish the earned-categorical-machine theorem **??** of the introduction.

Remark 7.17 (Planned Lean module `TauLib.BookII.Holomorphy.EarnedCat`, **[τ-Effective]**)**.** Each earned law becomes a Lean theorem (not axiom) in the planned module. The key proof techniques are:

- **Admissibility closure.** Typing and stability preservation under concatenation reduce to finite `decide`-able predicates on NF codes, given the existing `Decidable` instances for Typed, Stable, and tail-independence at a fixed depth.
- **Identity transformer.** Defined by primitive recursion as the empty list of elementary rewriting operations; unit laws are definitional after a `simp` with the list-append lemmas.

- **Associativity.** Reduces to the Church–Rosser theorem for the τ-native rewriting system, which is proved in `TauLib.BookI.Addressability` (Hinge 7's scope); here it is imported as a lemma and applied via `rfl` up to Norm.
- **Functoriality.** Compositional `simp`-lemmas reducing ρ^* to code concatenation; boundary-preserving closure is a structure field on the probe-category type.

No `sorry` is anticipated in this module. The dependencies are minimal: the admissibility-predicate Lemmas of `TauLib.BookII.Holomorphy.HolMaps` and the confluence theorem of `TauLib.BookI.Addressability`. The earned-categorical machinery is thus a clean downstream consequence of two independent pieces of infrastructure.

Remark 7.18 (Closing remark of §??)**.** With Corollary **??** in hand, the boundary-first programme of Hinges 4 and 5 has reached its categorical ground-floor: every piece of categorical structure that will be used downstream (in §§**??**–**??** of this paper, in Hinge 6's τ-topos, and in Hinge 7's canonical-addressability construction) has been derived from ω-germ transformers and three admissibility predicates. No category axiom is imposed; all are earned. This is the structural sense in which τ-holomorphy is the ontological primary: it *generates* the categorical apparatus rather than presupposing it. The remaining sections of this paper harvest the consequences.

8. σ-ANTI-HOLOMORPHY AND IDEMPOTENT-SUPPORTED HOLOMORPHY

8.1 The sigma-involution on carriers and scalar codomains

Hinge 4 established that the scalar codomain $\mathbb{D}$ carries a canonical involutive automorphism $\sigma\colon \mathbb{D} \to \mathbb{D}$ characterized by its action on the lemniscate idempotents:

$$\sigma(e_+) \;=\; e_-, \qquad \sigma(e_-) \;=\; e_+, \qquad \sigma(j) \;=\; -j, \qquad \sigma(1) \;=\; 1, \qquad \sigma(0) \;=\; 0. \tag{6}$$

Equation (**??**) determines σ uniquely as an $\mathcal{R}'_\partial$-algebra automorphism of $\mathbb{D}$: writing an element in its canonical form $z = z_+\,e_+ + z_-\,e_-$ with $z_\pm \in \mathcal{R}'_\partial$, one has

$$\sigma(z_+\,e_+ + z_-\,e_-) \;=\; z_-\,e_+ + z_+\,e_-. \tag{7}$$

This is the *lemniscate lobe-swap*: the two sectors $\mathcal{R}'_\partial \cdot e_+$ and $\mathcal{R}'_\partial \cdot e_-$ are exchanged, while the *real axis* $\mathcal{R}'_\partial \cdot 1 \subset \mathbb{D}$ (the scalar diagonal spanned by the unit $1 = e_+ + e_-$) is fixed pointwise. The split-complex unit $j = e_+ - e_-$ is σ-antifixed: $\sigma(j) = \sigma(e_+) - \sigma(e_-) = e_- - e_+ = -j$. One checks immediately that $\sigma^2 = \mathrm{id}_{\mathbb{D}}$ and that σ is an $\mathcal{R}'_\partial$-algebra homomorphism (it preserves $+, \cdot, 0, 1$ and the $\mathcal{R}'_\partial$-action on $\mathbb{D}$) [**?**].

The involution extends canonically to every admissible carrier. Let X be an object of the τ-transformer category, viewed as a carrier equipped with the Hinge 4 lemniscate structure $X = X_+ \vee X_-$. The σ-involution $\sigma_X\colon X \to X$ is the structural lobe-swap that exchanges X_+ and X_- while fixing the crossing-point germ. For the distinguished scalar codomain $Y = \mathbb{D}$, the involution $\sigma_{\mathbb{D}}$ coincides with σ as in (**??**).

Remark 8.1 (σ as a depth-0 τ-holomorphic map)**.** The lobe-swap $\sigma_X\colon X \to X$ is itself an admissible τ-holomorphic map: it is represented by a normal-form code of depth 0, carries the typing $\mathsf{Typed}(X, X, \sigma_X)$ tautologically, is τ-equality-stable (since it is an $\mathcal{R}'_\partial$-algebra involution and hence preserves every $\mathcal{R}'_\partial$-linear relation), and satisfies tail-independence trivially. Thus $\sigma_X \in \mathsf{Hol}_\tau(X, X)$ for every admissible carrier X. This is the pivotal fact that drives Theorem **??** below.

Given $f \in \mathsf{Hol}_\tau(X, Y)$ and the σ-involutions on both the domain and the codomain, define the *σ-conjugate transformer*

$$\bar{f} \;:=\; \sigma_Y \circ f \circ \sigma_X. \tag{8}$$

Intuitively, $\bar{f}$ is what one obtains by swapping lobes at the input, pushing through f, and swapping lobes again at the output. The remainder of this section establishes that (**??**) is the τ-native substitute for complex conjugation at the transformer level.

8.2 The sigma-anti-holomorphy theorem

Theorem 8.2 (σ-Anti-Holomorphy, **[τ-Effective]**). *For every admissible τ-holomorphic map $f \in \mathsf{Hol}_\tau(X, Y)$, the σ-conjugate*

$$\bar{f} \;=\; \sigma_Y \circ f \circ \sigma_X \tag{9}$$

is again an admissible τ-holomorphic map, i.e. $\bar{f} \in \mathsf{Hol}_\tau(X, Y)$. *The assignment* $f \mapsto \bar{f}$ *is an involution of* $\mathsf{Hol}_\tau(X, Y)$*:* $\bar{\bar{f}} = f$*. In particular,* anti-holomorphic *transformers in the classical sense are themselves τ-holomorphic — no separate calculus of* $\bar{z}$*-derivatives is required.*

Proof. We verify admissibility of $\bar{f}$ by reducing to Earned Composition (§**??**). By Remark **??**, $\sigma_X \in \mathsf{Hol}_\tau(X, X)$ and $\sigma_Y \in \mathsf{Hol}_\tau(Y, Y)$. The hypothesis gives $f \in \mathsf{Hol}_\tau(X, Y)$. Earned Composition states that Hol_τ is closed under composition: if $g \in \mathsf{Hol}_\tau(A, B)$ and $h \in \mathsf{Hol}_\tau(B, C)$ then $h \circ g \in \mathsf{Hol}_\tau(A, C)$. Applying this twice yields

$$f \circ \sigma_X \;\in\; \mathsf{Hol}_\tau(X, Y), \qquad \bar{f} \;=\; \sigma_Y \circ (f \circ \sigma_X) \;\in\; \mathsf{Hol}_\tau(X, Y),$$

which is the stated admissibility.

For the involutive property, compute using $\sigma_X \circ \sigma_X = \mathrm{id}_X$ and $\sigma_Y \circ \sigma_Y = \mathrm{id}_Y$:

$$\bar{\bar{f}} \;=\; \sigma_Y \circ \bar{f} \circ \sigma_X \;=\; \sigma_Y \circ (\sigma_Y \circ f \circ \sigma_X) \circ \sigma_X \;=\; (\sigma_Y \circ \sigma_Y) \circ f \circ (\sigma_X \circ \sigma_X) \;=\; f.$$ □

Remark 8.3 (The unified bipolar sense of anti-holomorphy). Theorem **??** marks a sharp departure from orthodox complex analysis. Over $\mathbb{C}$, the prototypical anti-holomorphic map $z \mapsto \bar{z}$ fails the Cauchy–Riemann equations; it satisfies the *conjugate* Cauchy–Riemann equations and generates a separate formal calculus (the Wirtinger $\partial/\partial\bar{z}$ machinery). In the τ-framework, by contrast, σ-conjugation is *internal to the carrier structure*: it is the lemniscate lobe-swap σ_X, itself a depth-0 admissible τ-holomorphic map. Composing with σ_X and σ_Y therefore never leaves Hol_τ, and what classical analysis views as a separate anti-holomorphic category is reabsorbed into the τ-holomorphic category under a canonical involution. This is the *unified bipolar sense* of τ-holomorphy: the two lemniscate branches are connected through σ, and σ-conjugation is itself a τ-holomorphic operation (Book I [**?**], anti-holomorphy-lobe-swap identification).

Remark 8.4 (σ-fixed and σ-antifixed transformers). Since $f \mapsto \bar{f}$ is an involution of the abelian group $\mathsf{Hol}_\tau(X, Y)$ (addition is pointwise and commutes with σ), every f admits the canonical splitting

$$f \;=\; f^{\mathrm{sym}} + f^{\mathrm{asym}}, \qquad f^{\mathrm{sym}} \;:=\; \tfrac{1}{2}(f + \bar{f}), \qquad f^{\mathrm{asym}} \;:=\; \tfrac{1}{2}(f - \bar{f}),$$

with $\overline{f^{\mathrm{sym}}} = f^{\mathrm{sym}}$ and $\overline{f^{\mathrm{asym}}} = -f^{\mathrm{asym}}$. The σ-symmetric component takes values in the σ-fixed *real axis* $\mathcal{R}'_\partial \cdot 1 \subset \mathbb{D}$ when $Y = \mathbb{D}$; the σ-antisymmetric component takes values in the σ-antifixed *imaginary axis* $\mathcal{R}'_\partial \cdot j = \mathcal{R}'_\partial \cdot (e_+ - e_-)$.

8.3 Setup for idempotent-supported holomorphy

We now lift the lemniscate idempotent structure from scalars to transformers. Recall from Hinge 4 the canonical $\mathcal{R}'_\partial$-algebra splitting

$$\mathbb{D} \;\cong\; \mathcal{R}'_\partial \cdot e_+ \;\oplus\; \mathcal{R}'_\partial \cdot e_-, \qquad z \;=\; z_+ e_+ + z_- e_-, \qquad z_\pm \in \mathcal{R}'_\partial, \tag{10}$$

with the idempotent relations $e_+ + e_- = 1$, $e_+ \cdot e_- = 0$, $e_+^2 = e_+$, $e_-^2 = e_-$, and the split-complex unit $j = e_+ - e_-$ [**?**, **?**]. The projections onto the two factors are

$$\pi_+ : \mathbb{D} \to \mathcal{R}'_\partial, \quad z = z_+ e_+ + z_- e_- \;\mapsto\; z_+, \qquad \pi_- : \mathbb{D} \to \mathcal{R}'_\partial, \quad z \;\mapsto\; z_-. \tag{11}$$

Equivalently, $\pi_+(z) = z \cdot e_+$ viewed in the e_+-summand of $\mathbb{D}$ and identified with its $\mathcal{R}'_\partial$-coordinate; similarly for π_-. Each $\pi_\pm$ is an $\mathcal{R}'_\partial$-linear, $\mathcal{R}'_\partial$-algebra homomorphism $\mathbb{D} \to \mathcal{R}'_\partial$, hence is τ-equality-stable and admits a normal-form code of depth 0. Therefore, by the same reasoning as in Remark **??**, each projection is itself an admissible τ-holomorphic map:

$$\pi_\pm \;\in\; \mathsf{Hol}_\tau(\mathbb{D}, \mathcal{R}'_\partial). \tag{12}$$

Given $f \in \mathsf{Hol}_\tau(X, \mathbb{D})$, define the *sector components*

$$f_+ \;:=\; \pi_+ \circ f, \qquad f_- \;:=\; \pi_- \circ f, \tag{13}$$

so that $f(x) = f_+(x)\, e_+ + f_-(x)\, e_-$ for every $x \in X$. This is the pointwise analogue of the scalar decomposition (**??**). Theorem **??** below asserts that the sector components $f_\pm$ are themselves admissible τ-holomorphic maps into the scalar ring $\mathcal{R}'_\partial$, and that this decomposition is unique.

8.4 The idempotent-supported holomorphy theorem

Theorem 8.5 (Idempotent-Supported Holomorphy, **[τ-Effective]**)**.** *Every admissible τ-holomorphic map into $\mathbb{D}$ factors through the two lemniscate lobes. Concretely,*

$$\mathsf{Hol}_\tau(X, \mathbb{D}) \;=\; e_+ \cdot \mathsf{Hol}_\tau(X, \mathcal{R}'_\partial) \;\oplus\; e_- \cdot \mathsf{Hol}_\tau(X, \mathcal{R}'_\partial), \tag{14}$$

where the direct sum is internal and the two summands are annihilated by opposite idempotents. Equivalently, for every $f \in \mathsf{Hol}_\tau(X, \mathbb{D})$ there exists a unique pair $(f_+, f_-) \in \mathsf{Hol}_\tau(X, \mathcal{R}'_\partial)^2$ with

$$f \;=\; e_+\, f_+ \;+\; e_-\, f_-. \tag{15}$$

The sector components $f_\pm$ are given by $f_\pm = \pi_\pm \circ f$.

Proof. *Existence.* Given $f \in \mathsf{Hol}_\tau(X, \mathbb{D})$ with admissible normal-form code c, let $c_\pm$ be the post-composition of c with the idempotent projection code of $\pi_\pm$. We verify the three admissibility conditions for $c_\pm$:

(i) *Typing* $\mathsf{Typed}(X, \mathcal{R}'_\partial, c_\pm)$. The projection $\pi_\pm : \mathbb{D} \to \mathcal{R}'_\partial$ is typed $\mathbb{D} \to \mathcal{R}'_\partial$ at depth 0 (it is an $\mathcal{R}'_\partial$-algebra homomorphism) and c is typed $X \to \mathbb{D}$ by hypothesis; the composite code is typed $X \to \mathcal{R}'_\partial$.

(ii) *τ-stability* $\mathsf{Stable}(X, \mathcal{R}'_\partial, c_\pm)$. Each $\pi_\pm$ is $\mathcal{R}'_\partial$-linear and in particular preserves $\sim$ (§**??**), so post-composition with $\pi_\pm$ preserves stability.

(iii) *Tail-independence.* Since $\pi_\pm$ has depth 0, it adds no new tails; the tail-independence of $c_\pm$ is inherited directly from that of c.

By Earned Composition (§**??**), the composite $\pi_\pm \circ f$ is again admissible, so $f_\pm := \pi_\pm \circ f \in \mathsf{Hol}_\tau(X, \mathcal{R}'_\partial)$. The identity (**??**) holds pointwise because $\mathbb{D}$ splits as in (**??**): for every $x \in X$,

$$f(x) \;=\; \pi_+(f(x))\, e_+ + \pi_-(f(x))\, e_- \;=\; e_+\, f_+(x) + e_-\, f_-(x).$$

Uniqueness. Suppose $f = e_+ f_+ + e_- f_- = e_+ f'_+ + e_- f'_-$ with $f_\pm, f'_\pm \in \mathsf{Hol}_\tau(X, \mathcal{R}'_\partial)$. Multiply both sides by e_+. Using $e_+^2 = e_+$ and $e_+ e_- = 0$,

$$e_+ f \;=\; e_+(e_+ f_+ + e_- f_-) \;=\; e_+ f_+,$$

and analogously $e_+ f = e_+ f'_+$. Hence $e_+ f_+ = e_+ f'_+$ in $\mathsf{Hol}_\tau(X, \mathbb{D})$; applying π_+ to both sides gives $f_+ = f'_+$ in $\mathsf{Hol}_\tau(X, \mathcal{R}'_\partial)$. The same argument with e_- yields $f_- = f'_-$.

Direct-sum decomposition. The existence and uniqueness statements together say that the map

$$\mathsf{Hol}_\tau(X, \mathcal{R}'_\partial) \oplus \mathsf{Hol}_\tau(X, \mathcal{R}'_\partial) \;\longrightarrow\; \mathsf{Hol}_\tau(X, \mathbb{D}), \qquad (g, h) \;\longmapsto\; e_+ g + e_- h$$

is bijective. The two summands $e_+ \cdot \mathsf{Hol}_\tau(X, \mathcal{R}'_\partial)$ and $e_- \cdot \mathsf{Hol}_\tau(X, \mathcal{R}'_\partial)$ intersect trivially (multiplying $e_+ g = e_- h$ by e_+ gives $e_+ g = 0$, hence $g = 0$, and similarly $h = 0$), so (**??**) is an internal direct sum of $\mathsf{Hol}_\tau(X, \mathcal{R}'_\partial)$-modules. ☐

Remark 8.6 (Restriction to lobes)**.** Equation (**??**) exhibits $\mathsf{Hol}_\tau(X, \mathbb{D})$ as a free $\mathsf{Hol}_\tau(X, \mathcal{R}'_\partial)$-module of rank 2, with canonical basis $\{e_+, e_-\}$. Multiplication by e_+ is the projection onto the first lobe, and each lobe $e_+ \cdot \mathsf{Hol}_\tau(X, \mathcal{R}'_\partial)$ is canonically isomorphic to $\mathsf{Hol}_\tau(X, \mathcal{R}'_\partial)$. The decomposition is therefore the transformer-level analogue of the scalar splitting (**??**) and is the structural content of the categorical identity $\mathbb{D} \cong \mathcal{R}'_\partial \cdot e_+ \oplus \mathcal{R}'_\partial \cdot e_-$ once lifted across the functor $\mathsf{Hol}_\tau(X, -)$.

8.5 Irreducibility and single-lobe support

Definition 8.7 (Irreducible transformer). *A map $f \in \mathsf{Hol}_\tau(X, \mathbb{D})$ is* irreducible *if it admits no nontrivial decomposition $f = f^{(1)} + f^{(2)}$ with both $f^{(1)}, f^{(2)} \in \mathsf{Hol}_\tau(X, \mathbb{D})$ nonzero and $f^{(1)} \cdot f^{(2)} = 0$ in $\mathsf{Hol}_\tau(X, \mathbb{D})$.*

The zero-product constraint picks out decompositions into $\mathcal{R}'_\partial$-algebra orthogonal summands; it is automatically satisfied by the lobe decomposition $f = e_+ f_+ + e_- f_-$ because $e_+ e_- = 0$ in $\mathbb{D}$.

Proposition 8.8 (Single-lobe support for irreducibles). *Let $f \in \mathsf{Hol}_\tau(X, \mathbb{D})$ be irreducible and nonzero. Then exactly one of $f_+, f_- \in \mathsf{Hol}_\tau(X, \mathcal{R}'_\partial)$ is nonzero. Equivalently, f is supported on exactly one lobe: either $f \in e_+ \cdot \mathsf{Hol}_\tau(X, \mathcal{R}'_\partial)$ or $f \in e_- \cdot \mathsf{Hol}_\tau(X, \mathcal{R}'_\partial)$.*

Proof. By Theorem **??** we have $f = e_+ f_+ + e_- f_-$ with $f_\pm \in \mathsf{Hol}_\tau(X, \mathcal{R}'_\partial)$. Setting $f^{(1)} := e_+ f_+$ and $f^{(2)} := e_- f_-$ gives a decomposition in $\mathsf{Hol}_\tau(X, \mathbb{D})$ with

$$f^{(1)} \cdot f^{(2)} \;=\; (e_+ f_+)(e_- f_-) \;=\; (e_+ e_-)\, f_+ f_- \;=\; 0.$$

By irreducibility one of $f^{(1)}, f^{(2)}$ must be zero, i.e. one of f_+, f_- is zero. Since $f \neq 0$, the other is nonzero. □

A general (possibly reducible) $f \in \mathsf{Hol}_\tau(X, \mathbb{D})$ then decomposes canonically into at most two orthogonal single-lobe components, one in each summand of (**??**). This recovers at the transformer level the Hinge 4 principle that *holomorphic transformers are supported in exactly one idempotent component per lobe* [**?**].

8.6 The sigma-equivariant endomorphism monoid

The σ-involution on carriers selects a distinguished sub-monoid of τ-holomorphic endomorphisms.

Definition 8.9 (σ-equivariant endomorphism). *A τ-holomorphic endomorphism $f\colon X \to X$ is* σ-equivariant *if $f \circ \sigma_X = \sigma_X \circ f$, equivalently if $\bar{f} = f$. Denote by*

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

the set of σ-equivariant τ-holomorphic endomorphisms of X.

Proposition 8.10 ($\mathsf{HolEnd}^\sigma_\tau(X)$ is a sub-monoid). *$\mathsf{HolEnd}^\sigma_\tau(X)$ contains id_X and is closed under composition; it is therefore a sub-monoid of $\mathsf{HolEnd}_\tau(X)$.*

Proof. The identity is σ-equivariant since $\mathrm{id}_X \circ \sigma_X = \sigma_X = \sigma_X \circ \mathrm{id}_X$. For $f, g \in \mathsf{HolEnd}^\sigma_\tau(X)$,

$$(f \circ g) \circ \sigma_X \;=\; f \circ (g \circ \sigma_X) \;=\; f \circ (\sigma_X \circ g) \;=\; (f \circ \sigma_X) \circ g \;=\; (\sigma_X \circ f) \circ g \;=\; \sigma_X \circ (f \circ g),$$

so $f \circ g \in \mathsf{HolEnd}^\sigma_\tau(X)$. Closure under composition within $\mathsf{HolEnd}_\tau(X)$ is Earned Composition. □

Structurally, $\mathsf{HolEnd}^\sigma_\tau(X)$ is the natural host for σ-fixed scalars and σ-symmetric constructions: both lobes of X are acted on symmetrically, and the fixed-point locus of σ — the crossing-point germ of the lemniscate — is preserved pointwise. Hinge 3's crossing-point germ uniqueness [**?**] may be understood as the statement that $\mathsf{HolEnd}^\sigma_\tau(X)$ has a distinguished fixed germ at which the σ-symmetric scalar ι_τ is pinned. The forward section §**??** returns to the finer structure of HolEnd_τ and its symmetric sub-monoid.

8.7 Consequences and corollaries

Corollary 8.11 (σ-fixed scalars lie on the real axis, **[τ-Effective]**). *If $s \in \mathbb{D}$ satisfies $\sigma(s) = s$, then $s = s_0 \cdot 1 = s_0 \in \mathcal{R}'_\partial$ for some $s_0 \in \mathcal{R}'_\partial$. Consequently, every σ-equivariant scalar-valued τ-holomorphic map $f \in \mathsf{Hol}_\tau(X, \mathbb{D})$ with $\bar{f} = f$ takes values in the* real axis *$\mathcal{R}'_\partial \cdot 1 \subset \mathbb{D}$, and is equivalently described by the underlying $\mathcal{R}'_\partial$-valued map $f_0 := \pi_+ \circ f = \pi_- \circ f \in \mathsf{Hol}_\tau(X, \mathcal{R}'_\partial)$.*

Proof. Write $s = s_+e_+ + s_-e_-$. The fixed-point condition (??) gives $s_-e_+ + s_+e_- = s_+e_+ + s_-e_-$, which forces $s_+ = s_-$. Setting $s_0 := s_+$ then yields $s = s_0(e_+ + e_-) = s_0 \cdot 1 = s_0 \in \mathcal{R}'_\partial$. Applying this pointwise to a σ-equivariant f, the identity $\bar{f} = f$ reads $\sigma(f(\sigma_X(x))) = f(x)$ for every $x \in X$; restricting to the fixed-point germ of σ_X and extending by the uniqueness statement of Theorem **??** gives $f_+ = f_-$. □

Corollary 8.12 (ι_τ as the unique balanced σ-fixed scalar, **[τ-Effective]**). *The master constant $\iota_\tau = 2/(\pi + e)$ of Hinge 3 [?] is σ-fixed and non-idempotent. It therefore lies on the real axis: $\iota_\tau \in \mathcal{R}'_\partial \cdot 1 \subset \mathbb{D}$, and is the unique balanced σ-symmetric scalar realising the Hinge 3 crossing-point geometry.*

Proof. σ-fixedness is a structural property of Hinge 3 (the lobe-swap fixes ι_τ pointwise because it sits at the lemniscate crossing). Non-idempotency $\iota_\tau^2 \neq \iota_\tau$ is Hinge 3, Proposition 1. Applying Corollary **??** places ι_τ in $\mathcal{R}'_\partial \cdot 1$. □

Corollary 8.13 (σ-action on sector decompositions, **[τ-Effective]**). *For every $f \in \mathsf{Hol}_\tau(X, \mathbb{D})$ with sector decomposition* (**??**),

$$\sigma(e_+f_+ + e_-f_-) \;=\; e_-f_+ + e_+f_-, \tag{16}$$

so that σ-conjugation on $\mathsf{Hol}_\tau(X, \mathbb{D})$ acts as the lobe-swap $(f_+, f_-) \mapsto (f_-, f_+)$ of the pair of sector components. In particular, f is σ-fixed if and only if $f_+ = f_-$, and in that case $f = f_+ \cdot 1 = f_+ \in \mathsf{Hol}_\tau(X, \mathcal{R}'_\partial)$, i.e. f is a real-valued τ-holomorphic map.

Proof. Apply σ pointwise and use $\sigma(e_+) = e_-$, $\sigma(e_-) = e_+$ together with $\sigma|_{\mathcal{R}'_\partial} = \mathrm{id}_{\mathcal{R}'_\partial}$: $\sigma(e_+f_+(x) + e_-f_-(x)) = e_-f_+(x) + e_+f_-(x)$. Rewriting with the pair (f_+, f_-) gives the claimed swap. □

Corollary **??** makes precise the intertwining of the σ-involution with the idempotent decomposition: σ and the lobe projections $\pi_\pm$ generate a faithful action of $\mathbb{Z}/2$ on the free $\mathsf{Hol}_\tau(X, \mathcal{R}'_\partial)$-module (**??**) by swap of generators, and the σ-fixed points are exactly the diagonal summand.

8.8 Registry and Lean preview

The constructions of this section are slated for the Lean module `TauLib.BookII.Holomorphy.SigmaIdem`, with Theorems **??** and **??** as the two headline lemmas. The sector decomposition (**??**) is foundational for the Book IV physics applications: the two lemniscate lobes e_+ and e_- correspond to the γ / η electromagnetic / strong force channels under the Hinge 4 dictionary [?, ?, ?], so single-lobe support (Proposition **??**) is the categorical form of charge-sector purity for irreducible τ-transformers. The σ-equivariant sub-monoid $\mathsf{HolEnd}^\sigma_\tau$ will be taken up again in §**??** as the carrier of the symmetric transformer calculus that drives Hinge 6's earned differentials.

9. THE HolEnd_τ CATEGORY VIA PRE-YONEDA COLLAPSE

9.1 Construction of HolEnd_τ

By the point §**??** has been established, every categorical constituent we need to assemble an endomorphism category has already been earned as theorem: carriers come with a probe category of primorial-depth refinements (§**??**); admissible ω-germ transformers form the morphism spaces $\mathsf{Hol}_\tau(X, Y)$ (§**??**); composition is defined by normalised sequential action on tails with code concatenation (Theorem **??**(a)); and a canonical tail-fixing NF code supplies the identity (Theorem **??**(b)). The holomorphic endomorphism category HolEnd_τ is the restriction of this earned structure to the diagonal case $Y = X$.

Definition 9.1 (The category HolEnd_τ). *The* holomorphic endomorphism category *of τ is the category with*

- Objects. *Pairs (X, f) with $X \in \mathrm{Obj}(\tau)$ a τ-admissible carrier and $f \in \mathsf{Hol}_\tau(X, X)$ an admissible τ-holomorphic endomorphism of X in the sense of Definition* **??**.
- Morphisms. *A morphism $\phi\colon (X, f) \to (Y, g)$ is an admissible transformer $\phi \in \mathsf{Hol}_\tau(X, Y)$ satisfying the intertwining condition*

$$g \circ \phi \;=\; \phi \circ f \qquad \textit{in } \mathsf{Hol}_\tau(X, Y),$$

*where composition is the earned composition of §**??** and equality is NF-equivalence of codes.*

- Identity. $\mathrm{id}_{(X,f)} := \mathrm{id}_X$*, which intertwines f with itself trivially.*

- Composition. *The composite of* $\phi\colon (X,f) \to (Y,g)$ *and* $\psi\colon (Y,g) \to (Z,h)$ *is the earned composition* $\psi \circ \phi \in \mathsf{Hol}_\tau(X,Z)$, *which is again an intertwiner between* f *and* h *by horizontal pasting of the two commuting squares.*

Proposition 9.2 (HolEnd$_\tau$ is a category **[τ-Effective]**). *The data of Definition* **??** *form a category: composition of intertwiners is an intertwiner, identity intertwines trivially, associativity and unit laws are inherited from the earned categorical machine of Theorem* **??**.

Lean-grade sketch. Closure under composition. Given ϕ intertwining (f,g) and ψ intertwining (g,h), the equality $h \circ (\psi \circ \phi) = (\psi \circ \phi) \circ f$ follows by two applications of earned associativity (Theorem **??**(c)) and the two local intertwining relations:

$$h \circ \psi \circ \phi \;=\; \psi \circ g \circ \phi \;=\; \psi \circ \phi \circ f.$$

Identity. $g \circ \mathrm{id}_X = g = \mathrm{id}_X \circ g$ by Theorem **??**(b). *Associativity and unit laws.* Inherited verbatim from Theorem **??**. □

Remark 9.3 (What has already been earned). No new axiom is introduced in Definition **??**. The entire content of the definition is *structural selection* within an already-earned larger category: we keep those objects that carry an admissible self-transformer, and we keep those morphisms that intertwine. The non-trivial content of Theorem **??** is not the categorical assembly (which is automatic, as just proved) but the additional claim that HolEnd$_\tau$ is *concretely* representable in the boundary addressable objects of Hinge 1 — the pre-Yoneda collapse of §**??** below.

9.2 The probe category P_τ (recap)

We briefly recall the probe category P_τ from §**??**.5, which provides the "small site" over which the pre-Yoneda collapse will be performed.

Definition 9.4 (Probe category P_τ). *The* probe category of τ-carriers *has as objects the primorial-depth carrier types* X_n *(one per admissible primorial stage) and as morphisms the reduction / refinement transformers* $X_n \to X_m$ *for* $n \geq m$ *that respect the tail discipline of Definition* **??**.

Remark 9.5 (Thin posetal structure). P_τ is *thin*: at most one morphism $X_n \to X_m$ exists for any ordered pair, because refinement at primorial depth is unique up to NF reduction. In particular, P_τ is posetal, and its opposite P_τ^{op} is again thin and posetal. Thinness will be essential below: it is what makes presheaves on P_τ^{op} fully determined by their value on the chain of primorial stages.

Remark 9.6 (P_τ is the site for pre-Yoneda). The pre-Yoneda collapse of §**??** embeds τ into $\mathbf{PSh}(P_\tau) := [\,P_\tau^{\mathrm{op}}, \mathbf{Set}\,]$ by the assignment $X \mapsto \mathsf{Hol}_\tau(-,X)$. We now explain why this particular assignment *collapses* at the presheaf level: every admissible carrier is represented by a canonical boundary-addressed code rather than by a genuine presheaf requiring Grothendieck descent.

9.3 The Yoneda embedding, classical picture

For a locally small category C, the Yoneda embedding is

$$y\colon C \longrightarrow \mathbf{PSh}(C), \qquad X \;\mapsto\; y_X := \mathrm{Hom}_C(-,X),$$

a functor $C^{\mathrm{op}} \to \mathbf{Set}$ assigning to each object X its representable presheaf. The classical Yoneda lemma states that y is fully faithful: $\mathrm{Hom}_{\mathbf{PSh}(C)}(y_X, y_Y) \cong \mathrm{Hom}_C(X,Y)$ naturally in X and Y. However, y is generally *not* essentially surjective: most presheaves on C are not representable. The representable ones form a thin subcategory of $\mathbf{PSh}(C)$.

In the τ-setting, the candidate embedding is

$$y_\tau\colon \tau \longrightarrow \mathbf{PSh}(P_\tau), \qquad X \;\mapsto\; \mathsf{Hol}_\tau(-,X).$$

A priori, $\mathsf{Hol}_\tau(-,X)$ is a presheaf on the probe category P_τ^{op}: it assigns to each probe carrier X_n the countable set $\mathsf{Hol}_\tau(X_n, X)$ of admissible transformers, and to each refinement $X_n \to X_m$ the pullback of transformers. The question is whether this presheaf is representable by a canonical τ-object — and, if so, by what kind of object.

Classical Yoneda would answer: representable exactly when of the form $\mathrm{Hom}(-,X)$ for some $X \in C$. The τ-native answer, which we now establish, is stronger: *every* admissible X yields a representable presheaf, and the representative is a canonical boundary-addressed code living in $\partial\tau^3$ — the hyperfactorization boundary of Hinge 1 [**?**].

9.4 The pre-Yoneda collapse theorem

We now state the core structural result.

Theorem 9.7 (Pre-Yoneda collapse **[τ-Effective]**, modulo forthcoming Hinge 7 canonical-address NF confluence and Hinge 1 boundary-Hom structure)**.** *Under the τ-native admissibility predicates of §??, the Yoneda functor*

$$y_\tau\colon \tau \longrightarrow \mathbf{PSh}(P_\tau), \qquad X \mapsto \mathsf{Hol}_\tau(-, X),$$

collapses at the presheaf level*: for every admissible carrier $X \in \mathrm{Obj}(\tau)$, the presheaf $\mathsf{Hol}_\tau(-, X)$ is representable by a canonical code object*

$$c_X \in \partial\tau^3,$$

where $\partial\tau^3$ denotes the boundary of the hyperfactorization fibered product of Hinge 1 [?]. The assignment $X \mapsto c_X$ is well-defined up to canonical NF equivalence and functorial in P_τ-refinements.

The name *pre-Yoneda collapse* records three things at once: (i) the embedding happens *pre-categorically* (the codes c_X live on the boundary algebra of Hinge 4, not inside an a-priori-assumed category of presheaves); (ii) the codomain $\mathbf{PSh}(P_\tau)$ *collapses* to the thin image of $\partial\tau^3$, because every $\mathsf{Hol}_\tau(-, X)$ is representable by a code rather than being a genuine presheaf requiring descent data; and (iii) the collapse precedes the usual functor-of-points machinery, so no un-earned presheaf machinery is smuggled in.

Proof sketch. The argument has four steps, each using only machinery already earned in preceding sections or imported from Hinges 1–4.

Step 1: canonical ω-tail generator code. For each admissible carrier X, the tail-stability predicate Stable (Definition **??**) and tail-independence beyond a finite witness depth (Definition **??**) together guarantee that $\mathsf{Hol}_\tau(-, X)$ has a finite-witness generating set: any two admissible transformers $X_n \to X$ that agree on prefixes up to depth k_0 are equivalent under $\sim$. Choose the NF representative c_X of the canonical tail-fixing code of X itself (the code whose semantics is the inclusion of the X-tail into the ambient Ω_{tail}); this is uniquely determined by the NF reduction system of §?? (forthcoming Hinge 7 confluence theorem ensures strong normalisation, but for the present paper we only need weak NF on codes already in admissible form).

Step 2: addressability in $\partial\tau^3$. The boundary of the hyperfactorization fibered product $\partial\tau^3 = \partial(\tau^1 \times_f T^2)$ of Hinge 1 [?] is, by construction, the universal recipient of ABCD-addressed codes: each admissible code c sits on a canonical address $(a, b, c_{\text{coord}}, d) \in \partial\tau^3$ determined by its tower-atom decomposition. The code c_X inherits such an address from the primorial-depth stage at which X stabilises. Write $\mathrm{addr}(X) := c_X \in \partial\tau^3$.

Step 3: representability. We claim that the presheaf $\mathsf{Hol}_\tau(-, X)$ is represented by $c_X \in \partial\tau^3$. For each probe carrier X_n, the set $\mathsf{Hol}_\tau(X_n, X)$ is in canonical bijection with the set of admissible addresses from X_n to $\mathrm{addr}(X)$ in $\partial\tau^3$, i.e. with $\mathrm{Hom}_{\partial\tau^3}(X_n, c_X)$ interpreted in the boundary algebra. The bijection sends a transformer $\phi \in \mathsf{Hol}_\tau(X_n, X)$ to its NF code viewed as an address, and is natural in X_n (i.e. commutes with refinements in P_τ) because refinements are themselves transformers and admissibility is preserved under composition (Theorem **??**(a)).

Step 4: functoriality and uniqueness. The assignment $X \mapsto c_X$ is functorial on P_τ-refinements by the same naturality computation, and uniqueness of c_X up to canonical NF equivalence follows from weak confluence of the NF reduction system (Theorem **??**(c)). Hence y_τ factors through the thin image of $\partial\tau^3$ inside $\mathbf{PSh}(P_\tau)$, and this factorisation is canonical. □

Remark 9.8 (Why "collapse")**.** In a classical category, $y(C) \subsetneq \mathbf{PSh}(C)$ strictly: most presheaves are not representable. In the τ-setting, by contrast, the admissibility predicates built into Hol_τ are strong enough that y_τ is *essentially surjective onto a canonical boundary-addressed subcategory* — the image sits inside $\partial\tau^3$ as a thin subcategory rather than scattered across an ambient presheaf category. This is the "collapse": the normally open gap between representables and arbitrary presheaves is closed by τ's boundary structure. No descent data is required beyond the hyperfactorization address.

Remark 9.9 (Relation to Hinge 1)**.** The boundary object $\partial\tau^3$ is where Hinge 1's hyperfactorization structure [?] meets the categorical machinery of the present paper. Its concrete content is the set of tower-atom addresses $(A \uparrow\uparrow C)^B \cdot D$ for admissible ABCD coordinates; the pre-Yoneda collapse embeds each admissible carrier as one such address. This is the technical statement behind the heuristic slogan *"objects are addresses, not generators"*, which will be developed fully in Hinge 7.

9.5 HolEnd_τ as a concrete boundary-addressed category

The pre-Yoneda collapse of §?? makes HolEnd_τ *concrete*: its objects are not abstract pairs but boundary-addressed codes, and its morphisms are not abstract intertwiners but admissible boundary transformers preserving the intertwining relation.

Proposition 9.10 (HolEnd$_\tau$ is concretely representable **[τ-Effective]**). *Under the pre-Yoneda collapse of Theorem ??,*

$$\mathsf{HolEnd}_\tau \;\cong\; \{\, (c_X, f_X) \,\big|\, c_X \in \partial\tau^3,\; f_X \in \mathrm{End}_\tau(c_X) \,\},$$

where $\mathrm{End}_\tau(c_X) \subset \partial\tau^3$ *is the set of admissible boundary endomorphism codes at address* c_X*, and morphisms* $(c_X, f_X) \to (c_Y, f_Y)$ *are admissible boundary transformers* $\phi \in \mathsf{Hol}_\tau(c_X, c_Y)$ *with* $f_Y \circ \phi = \phi \circ f_X$ *in the NF code algebra.*

Sketch. Apply Theorem ?? to objects and morphisms separately. Objects: each $(X, f) \in \mathrm{Obj}(\mathsf{HolEnd}_\tau)$ maps to (c_X, f_X) where f_X is the NF code of f viewed as a boundary self-address. Morphisms: each intertwiner ϕ has an NF code, and the intertwining relation $g \circ \phi = \phi \circ f$ transports to its boundary-addressed shadow by functoriality of the pre-Yoneda collapse. Well-definedness up to NF equivalence follows from Theorem ??(c). □

Definition 9.11 (Automorphism sub-monoid). *For each* $(X, f) \in \mathrm{Obj}(\mathsf{HolEnd}_\tau)$*, the* automorphism group of (X, f) *is the sub-monoid of self-intertwiners that are invertible in* $\mathsf{Hol}_\tau(X, X)$*:*

$$\mathrm{Aut}_\tau(X, f) \;:=\; \{\, \phi \in \mathsf{HolEnd}_\tau((X, f), (X, f)) \,\big|\, \phi^{-1} \in \mathsf{Hol}_\tau(X, X) \,\}.$$

When $f = \mathrm{id}_X$*, we write* $\mathrm{Aut}_\tau(X) := \mathrm{Aut}_\tau(X, \mathrm{id}_X)$*, the* bare automorphism group of X.

Remark 9.12 (Invertibility is decidable). Invertibility in $\mathsf{Hol}_\tau(X, X)$ is decidable by NF witness: one searches the finite prefix depth k_0 of tail-independence for a two-sided inverse NF code; if none exists up to depth k_0, none exists at all. Hence $\mathrm{Aut}_\tau(X) \subset \mathsf{HolEnd}_\tau(X)$ is a decidable submonoid, and HolEnd_τ has well-defined (concrete, decidable) automorphism structure.

9.6 The σ-equivariant endomorphism monoid $\mathsf{HolEnd}_\tau^\sigma$

We now pass from HolEnd_τ to its σ-equivariant refinement. Recall from §?? that the σ-involution on the lemniscate induces an involution $\sigma\colon \mathsf{Hol}_\tau(X, X) \to \mathsf{Hol}_\tau(X, X)$ by $\sigma(f) := \sigma \circ f \circ \sigma$ (on the boundary algebra $\mathbb{D}$, σ acts as $e_+ \leftrightarrow e_-$).

Definition 9.13 (HolEnd$_\tau^\sigma$). *The* σ-equivariant holomorphic endomorphism monoid *is*

$$\mathsf{HolEnd}_\tau^\sigma \;:=\; \{\, (X, f) \in \mathrm{Obj}(\mathsf{HolEnd}_\tau) \,\big|\, \sigma(f) = f \,\},$$

with morphisms inherited from HolEnd_τ *restricted to σ-equivariant intertwiners ϕ satisfying $\sigma(\phi) = \phi$ as NF codes.*

Proposition 9.14 (HolEnd$_\tau^\sigma$ is a (wide) submonoidal subcategory **[τ-Effective]**). $\mathsf{HolEnd}_\tau^\sigma$ *is closed under the identity and composition of* HolEnd_τ*: if $\sigma(f) = f$ and $\sigma(g) = g$, then $\sigma(g \circ f) = \sigma(g) \circ \sigma(f) = g \circ f$ (using Theorem ??(a) and the fact that σ is a covariant involutive $\mathcal{R}'_\partial$-algebra automorphism on ω-germ transformers); and $\sigma(\mathrm{id}_X) = \mathrm{id}_X$ because σ fixes the canonical tail-fixing code of every carrier. Hence* $\mathsf{HolEnd}_\tau^\sigma \subseteq \mathsf{HolEnd}_\tau$ *is a wide subcategory (same objects allowed, fewer morphisms) and, restricted to a single object (X, f), forms a sub-monoid of* $\mathsf{HolEnd}_\tau(X, f)$.

Remark 9.15 (Why "monoid" rather than "category"). On a single object (X, f), the endomorphism hom $\mathsf{HolEnd}_\tau^\sigma(X, f) := \mathsf{HolEnd}_\tau^\sigma((X, f), (X, f))$ is a monoid under composition with id_X as unit. The category $\mathsf{HolEnd}_\tau^\sigma$ is the disjoint union of these monoids glued by σ-equivariant intertwiners; for purposes of the universality statement in §?? below, it is the monoidal structure on hom that matters, hence the name "σ-equivariant endomorphism monoid" in informal usage.

Remark 9.16 (Natural host for σ-fixed structure). Any quantity canonically derived from τ-holomorphy and invariant under σ must be expressible within $\mathsf{HolEnd}_\tau^\sigma$. In particular: the σ-fixed crossing germ of Hinge 3 [?] must live in $\mathsf{HolEnd}_\tau^\sigma$, and the only candidate scalar it can equal is a σ-fixed unit-ball element of the boundary algebra $\mathbb{D}$. This sets up §??.

9.7 The universality of ι_τ in $\mathsf{HolEnd}_\tau^\sigma$

The master constant $\iota_\tau = 2/(\pi + e)$ was derived in Hinge 3 [?] as the unique σ-fixed crossing-germ scalar on the lemniscate. We can now state this universality categorically: ι_τ is the unique σ-fixed, non-idempotent, balanced scalar in the unit ball of $\mathbb{D}$, and it is invariant under every element of $\mathsf{HolEnd}_\tau^\sigma$.

Theorem 9.17 (ι_τ as the universal σ-fixed scalar of $\mathsf{HolEnd}_\tau^\sigma$ **[τ-Effective]**)**.** *The master constant $\iota_\tau \in \mathbb{D}$ satisfies:*

(i) σ-fixed: $\sigma(\iota_\tau) = \iota_\tau$*;*

(ii) Non-idempotent: $\iota_\tau \notin \{0, e_+, e_-, 1\}$*;*

(iii) Balanced (on the real axis): $\iota_\tau = \iota_\tau \cdot (e_+ + e_-) = \iota_\tau \cdot 1 \in \mathcal{R}'_\partial \cdot 1 \subset \mathbb{D}$, *i.e. ι_τ lies on the σ-fixed real axis of $\mathbb{D}$ (the unit-diagonal spanned by $1 = e_+ + e_-$);*

(iv) Preserved by every σ-equivariant endomorphism: *for every endomorphism $g \in \mathsf{HolEnd}_\tau(\mathbb{D})$ that is σ-equivariant ($\sigma \circ g = g \circ \sigma$) and admissible as a τ-holomorphic transformer, the σ-fixed subspace $\mathbb{D}^\sigma = \mathcal{R}'_\partial \cdot 1$ is preserved by g, and ι_τ is invariant under g by the Hinge-3 uniqueness characterisation. (Here we use the transformer-native form "g preserves ι_τ" rather than the interior evaluation-at-a-point notation; the two are equivalent inside the σ-fixed scalar fibre $\mathcal{R}'_\partial \cdot 1$.)*

Moreover, ι_τ is uniquely characterised in $\mathbb{D}$ by conditions (i)–(iii) *together with the Hinge 3 crossing-germ normalisation, and condition* (iv) *then follows as a theorem. In categorical language, ι_τ is the σ-fixed universal scalar of* $\mathsf{HolEnd}_\tau^\sigma$.

Proof sketch. Conditions (i)–(iii) are the content of Hinge 3 [?]: the σ-fixed crossing germ is unique, is non-idempotent (the idempotents $e_\pm$ are excluded by the crossing-germ condition), and sits on the diagonal of the idempotent decomposition of $\mathbb{D}$ by the explicit formula $\iota_\tau = 2/(\pi + e)$ applied in the balanced basis $e_\pm = (1 \pm j)/2$.

For (iv): given $(\mathbb{D}, g) \in \mathsf{Obj}(\mathsf{HolEnd}_\tau^\sigma)$, the σ-equivariance condition $\sigma \circ g = g \circ \sigma$ forces the action of g to preserve the σ-fixed subalgebra $\mathbb{D}^\sigma = \mathcal{R}'_\partial \cdot 1$ (Corollary **??**). The restriction $g|_{\mathbb{D}^\sigma}$ is therefore a τ-holomorphic endomorphism of the one-dimensional $\mathcal{R}'_\partial$-algebra $\mathcal{R}'_\partial \cdot 1$. We unpack the reduction step in three clauses:

- *σ-fixedness inherited.* Every element of the image $g(\mathcal{R}'_\partial \cdot 1)$ is itself σ-fixed, because g commutes with σ and the source is already σ-fixed. In particular, if g transports ι_τ to any element $g \bullet \iota_\tau$ of $\mathbb{D}$, that image lies on the real axis $\mathcal{R}'_\partial \cdot 1$.
- *Balance inherited.* Idempotent decomposition in $\mathbb{D}$ is preserved by any $\mathcal{R}'_\partial$-algebra endomorphism of the σ-fixed subalgebra: $e_\pm \mapsto e_\pm$ (since $e_+ + e_- = 1$ is the unique idempotent sum landing on the real axis modulo swap, which is ruled out by σ-equivariance). Hence $g \bullet \iota_\tau$ is still *balanced* in the sense $g \bullet \iota_\tau = (g \bullet \iota_\tau) \cdot e_+ + (g \bullet \iota_\tau) \cdot e_-$ with equal lobe components, matching clause (iii).
- *Non-idempotence preserved.* If $g \bullet \iota_\tau$ were idempotent, then so would be its σ-fixed preimage; but ι_τ is non-idempotent by clause (ii), and $g|_{\mathcal{R}'_\partial \cdot 1}$ is an algebra endomorphism (hence idempotence-reflecting on the one-dimensional fibre).

Combining the three clauses, $g \bullet \iota_\tau$ is σ-fixed, balanced, and non-idempotent — the exact three conditions (i)–(iii) that uniquely characterise ι_τ by the Hinge-3 crossing-germ normalisation [?, Theorem 1.3]. Therefore g restricts to the identity on the Hinge-3 crossing-germ scalar $\iota_\tau \in \mathcal{R}'_\partial \cdot 1$. No interior evaluation-at-a-point is used: the argument is carried entirely by the commuting square $\sigma \circ g = g \circ \sigma$, the three clauses above, and Hinge-3 uniqueness.

Finally, uniqueness: any scalar $\lambda \in \mathbb{D}$ satisfying (i)–(iii) must coincide with the Hinge 3 crossing-germ scalar by the normalisation fixed there; and the crossing-germ scalar is $\iota_\tau = 2/(\pi + e)$ by [?, Theorem 1.3]. □

Remark 9.18 (What "universal" means here)**.** The sense of "universal" in Theorem **??** is not the Kan-extension universality of colimit objects; rather, it is the "unique fixed point of a canonical endomorphism monoid" universality familiar from Lawvere-style fixed-point theorems. Concretely: ι_τ is the unique element of the unit ball of $\mathbb{D}$ that is fixed by every σ-equivariant boundary endomorphism — there is no other σ-fixed, non-idempotent, balanced scalar. The role of $\mathsf{HolEnd}_\tau^\sigma$ is to provide the *categorical host* in which this universality can be stated precisely: the absence of $\mathsf{HolEnd}_\tau^\sigma$ would force us to restate Hinge 3's uniqueness theorem as an ad-hoc scalar identity, rather than as a canonical fixed-point statement.

9.8 Connection to $\mathbf{Cat}_\tau$ and forward outlook

The holomorphic endomorphism category HolEnd_τ and its σ-equivariant refinement $\mathsf{HolEnd}_\tau^\sigma$ are the *holomorphic cores* of the larger τ-topos $\mathbf{Cat}_\tau$ to be constructed in Hinge 6. Three forward links are worth flagging explicitly.

Pre-Yoneda collapse as the key technical input to the τ-topos. The subobject classifier Ω_τ of the τ-topos is internalised over the idempotent sublattice $B_\sigma(\mathbb{D}) = \{0, e_+, e_-, 1\} \subset \mathbb{D}$ (the four truth values of the Truth4 logic, paraconsistent in the sense of [?, ?]). The pre-Yoneda collapse theorem (Theorem **??**) is what allows this internalisation to succeed: because every admissible carrier factors canonically through $\partial\tau^3$, the subobject classifier can be defined pointwise on boundary addresses rather than as an abstract presheaf construction — in particular, no appeal is made here to the classical Grothendieck–Giraud topos machinery of [?, ?, ?, ?], though Hinge 6 will make explicit the precise relation between $\mathbf{Cat}_\tau$ and the classical notion of an elementary topos.

Boundary-addressed character of HolEnd_τ *feeds Hinge 7.* The concrete-representability result (Proposition **??**) places every object of HolEnd_τ at a canonical address in $\partial\tau^3$. The forthcoming Hinge 7 paper (*Address Resolution, Not Calculation*) lifts this addressing discipline to the genealogical DAG, the Cayley word metric, and the ontic ultrametric; the pre-Yoneda collapse is the categorical backbone of that lift.

The full τ-topos. $\mathbf{Hol}_\tau$ (from §**??**) and HolEnd_τ together form the locally small subcategory of the eventual τ-topos that inherits all of classical holomorphy's structural virtues (rigidity, boundary-determinacy, composition closure) while retaining the pre-categorical discipline of the present paper (no Cartesian products, no free diagonals, no unearned function spaces). The τ-topos of Hinge 6 is built on top of HolEnd_τ by adjoining the Truth4 internal logic; HolEnd_τ itself is *not* a topos, but its σ-equivariant refinement $\mathsf{HolEnd}_\tau^\sigma$ is the natural host of the ι_τ-universality that drives physical applications in Books IV–VII.

9.9 Registry and Lean preview

Remark 9.19 (Planned Lean module)**.** The results of this section will be formalised in `TauLib.BookII.Holomorphy.HolEnd`, with the principal artefacts:

- `HolEnd.lean` — the category HolEnd_τ of Definition **??**, with Proposition **??** as `HolEnd.is_category`.
- `PreYonedaCollapse.lean` — Theorem **??** via the canonical NF code assignment $X \mapsto c_X$; relies on the forthcoming NF confluence theorem from Hinge 7's planned `TauLib.BookI.Addressability` module.
- `HolEndConcrete.lean` — Proposition **??**, the concrete representability of HolEnd_τ in $\partial\tau^3$.
- `HolEndSigma.lean` — $\mathsf{HolEnd}_\tau^\sigma$ and Proposition **??**.
- `IotaUniversal.lean` — Theorem **??**, the universality of ι_τ as the σ-fixed $\mathsf{HolEnd}_\tau^\sigma$-invariant scalar; imports from `TauLib.BookIII.IotaTau.CrossingGerm` (Hinge 3 core).

Formal dependencies: from `TauLib.BookII.Holomorphy.EarnedCat` (Theorem **??**), `TauLib.BookII.Holomorphy.SigmaIdem` (Theorem **??**), `TauLib.BookIII.BoundaryAlgebra` (Hinge 4, [?]), and `TauLib.BookI.Hyperfactorization` (Hinge 1, [?]). The pre-Yoneda collapse is Lean-certifiable once the NF confluence theorem from Hinge 7 is discharged.

Remark 9.20 (Registry placeholder)**.** Registry ID II.T64 (HolEnd_τ *via pre-Yoneda collapse*) will be populated with Theorem **??**, Proposition **??**, and Theorem **??** upon v1 stabilisation of the paper bundle; provisional sub-IDs II.T64a–c track the three statements individually, mirroring the Hinge 4 sub-ID convention (cf. [?, ?, ?]).

10. FORWARD OUTLOOK: INTERIOR, HARTOGS, AND τ-NAVIER–STOKES

The present paper stops at the boundary algebra. Three natural extensions are handed off cleanly to later work.

Interior points and Hartogs continuation.. The boundary-first development here exhibits $\mathsf{Hol}_\tau(X, Y)$ and HolEnd_τ on the boundary. The extension to interior points — defined as stabilised tail-classes $\mathsf{Pt}_\tau(X) := \mathsf{Tail}_X/\sim$ — and to the Hartogs continuation operator $\mathsf{Hartogs}\colon \mathsf{Germ}_\tau(X) \to \mathsf{Int}_\tau(X)$ producing the unique interior section of an admissible boundary germ is the subject of a forthcoming Book II paper (*Global Hartogs without Local Infinity*). The key statement there is: every admissible boundary germ has a unique interior continuation, which is a normalisation operator rather than an axiom. The machinery developed here (admissibility, Stable, tail-independence, and the earned categorical apparatus) supplies all the groundwork that paper requires.

Topology-geometry parallel readouts.. The canonical ultrametric on ω-tails and the Archimedean incidence geometry of the lemniscate boundary are both *parallel readouts* of τ-holomorphy, glued by ω-coherence rather than by a metric-induces-topology argument. Establishing this inversion (topology earned from holomorphy, not the other way around) is the subject of the chapter on "holomorphy-to-geometry inversion" in the forthcoming Book II.

τ-Navier–Stokes as local Hartogs continuation.. Once interior-point construction is in place, the wave-equation signature of τ-holomorphy (Theorem **??**) opens a local-Hartogs route to the Navier–Stokes existence and regularity problem. The τ-Navier–Stokes programme recasts NS as admissible boundary data plus unique Hartogs continuation of a defect-functional; the hyperbolic PDE signature is natively built in. This programme belongs to Book III [**?**].

Riemann-Mapping analogues.. The boundary automorphism group $\mathrm{Aut}_\sigma(\mathbb{L})$ of the lemniscate plays a role analogous to the Möbius group for the disk. A Riemann-Mapping-Theorem-analogue for lemniscate domains, using the pre-Yoneda collapse developed here, is a natural subject for Book II.

11. LEAN ROADMAP AND ARTEFACTS

- `TauLib.BookII.Holomorphy.Tails` — ω-tails, prefix predicates, $\sim$-equivalence.
- `TauLib.BookII.Holomorphy.Germs` — admissible boundary germ type Germ_τ and decidability of equality.
- `TauLib.BookII.Holomorphy.HolMaps` — $\mathsf{Hol}_\tau(X, Y)$ as certified transformer type.
- `TauLib.BookII.Holomorphy.EarnedScalar` — scalar codomain uniqueness (forcing $\mathbb{D}$).
- `TauLib.BookII.Holomorphy.WaveCR` — the wave-equation theorem.
- `TauLib.BookII.Holomorphy.DiagonalDiscipline` — the no-Cartesian-product obstruction.
- `TauLib.BookII.Holomorphy.EarnedCat` — earned composition, identity, associativity, and functoriality.
- `TauLib.BookII.Holomorphy.SigmaIdem` — anti-holomorphy and idempotent-supported holomorphy.
- `TauLib.BookII.Holomorphy.HolEnd` — the HolEnd_τ category via pre-Yoneda collapse.

Formal dependencies: from Hinge 4's `TauLib.BookIII.BoundaryAlgebra` (see Book III registry IDs III.T81–T89). Forward dependencies: supplies the earned categorical machine to Hinge 6's `TauLib.BookI.TauTopos` and supplies transformer-addressing to Hinge 7's `TauLib.BookI.Addressability`.

12. REGISTRY IDENTIFIERS

Remark 12.1 (Registry IDs **[τ-Effective]**)**.** The eight main theorems of this paper are registered in `registry/book2_registry.jsonl` as II.T57 (admissible germs), II.T58 (holomorphy before mappings), II.T59 (earned scalar codomain), II.T60 (wave-equation Cauchy–Riemann), II.T61 (diagonal discipline), II.T62 (earned categorical machine), II.T63 (σ-anti-holomorphy and idempotent-supported holomorphy), II.T64 (HolEnd_τ via pre-Yoneda collapse), alongside the hinge-integration tabulation Theorem **??** (II.T65). Registry entries were added on 2026-04-22 following the v3.1 confirmation-panel CERTIFY verdict.

13. CONCLUSION AND FORWARD LINKS

We have established τ-holomorphy as the ontological primary of Category τ, defined directly as certified ω-germ transformers without prior notion of mapping, function, tuple, or Cartesian product. From this primary definition the entire categorical apparatus emerges as earned theorem: scalar codomain $\mathbb{D}$ is forced (Theorem **??**); the Cauchy–Riemann equations decouple into the hyperbolic wave equation (Theorem **??**); the diagonal-discipline theorem precludes elliptic collapse (Theorem **??**); and composition, identity, associativity, functoriality are derived rather than imposed (Theorem **??**).

The scope is deliberately cut at the boundary algebra — Hartogs continuation, interior points, topology-geometry inversion, and τ-Navier–Stokes are all handed to forthcoming Book II and Book III papers. What this paper supplies is the pre-categorical ontological root on which all those further developments will rest.

The forward architecture of the seven-hinge foundation bundle is:

- **Hinge 6** (forthcoming) — the τ-topos with Truth4 internal logic, using the earned categorical machine of §**??** and the four-atom dictionary of Hinge 4 as subobject classifier.
- **Hinge 7** (forthcoming) — canonical addressability, the genealogical DAG, and the ontic ultrametric, closing the loop back to Hinge 1's hyperfactorization.
- **Book II** [**?**] — Hartogs continuation, interior-point construction, holomorphy-to-geometry inversion.

- **Book III** [?] — τ-Navier–Stokes via local Hartogs continuation; the hyperbolic PDE signature established here is natively the right PDE type.
- **Books IV–VII** [?, ?, ?, ?] — sector-level physical couplings, quantum mechanics, general relativity, biological and metaphysical applications.

Acknowledgments

The authors thank the broader 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 1–4 whose structural objects this paper now lifts into a pre-categorical ontological primary.

Data and code availability

The source repository for the paper bundle is at `https://panta-rhei.site/papers/holomorphy-first`. Planned Lean 4 artefacts for the main theorems will appear in `TauLib.BookII.Holomorphy` (see §**??**).