HINGE 1 PREPRINT · V1.0 · APRIL 2026

The Hyperfactorization Theorem

Unique tower-atom decomposition in ZFC and Category τ

Thorsten Fuchs • Anna-Sophie Fuchs

Correspondence: thorsten@panta-rhei.site

April 2026

DOI: 10.5281/zenodo.19818957

ABSTRACT

We introduce a new normal form that refines the Fundamental Theorem of Arithmetic by recording tetration structure: every integer $X \geq 2$ admits a unique decomposition

$$X = (A \uparrow\uparrow C)^B \cdot D,$$

where A is the largest prime divisor of X; $C \geq 1$ is the maximal tetration height of the A-tower factor in X; $B \geq 1$ encodes the exponent of the resulting tower atom; and $D \geq 1$ has all prime factors strictly less than A. The associated *ABCD chart* $\Phi : \mathbb{N}_{\geq 2} \to \mathbb{P} \times \mathbb{N}^3_{\geq 1}$ is injective, so the chart is a complete invariant of the positive integers at or above 2. Uniqueness is proved via a rigidity result on the tower function $T(A, B, C) = (A \uparrow\uparrow C)^B$ — the *No-Tie Lemma* (adapted from Book I Chapter 22 [?]) — which shows that, under the maximality convention on C, the pair (B, C) is uniquely determined by the A-adic valuation of T. The maximality convention is essential: without it, three distinct tuples satisfy the bare equation at $X = 16$ (Example **??**(a)), and the theorem is literally false. The Fundamental Theorem of Arithmetic emerges as the height-$C = 1$ specialisation of the ABCD chart (Theorem **??**).

The proof (Theorem **??**) uses strong induction on X in four tightly-constrained steps: the Fundamental Theorem of Arithmetic fixes A; the A-adic valuation fixes the tower-atom valuation; the No-Tie Lemma fixes (B, C); cancellation fixes D. Every step is elementary and provable in Peano arithmetic (Appendix **??**).

Structural context.. The tower-atom decomposition arises natively inside Category τ as the *greedy peel algorithm* on τ-Idx (Theorem **??**, Book I Chapter 24 [?]), and the classical and categorical derivations agree pointwise under the kernel identification $\tau\text{-Idx} \cong \mathbb{N}_{\geq 1}$ (*Isomorphism Theorem*, Theorem **??**). The theorem is thus the first structural hinge of the Panta Rhei research programme, from which subsequent results — including the Prime Polarity Theorem [?] (Hinge 2), the Master Constant ι_τ paper [?] (Hinge 3), the Split-Complex Boundary Algebra paper [?] (Hinge 4), and the spectral-algebra constructions of Book III — derive. The ABCD chart's injectivity (Theorem **??**) implies the collapse of shadow equality to ontic identity in Category τ (Theorem **??**).

Lean status.. Computable checks (`hyperfact_check`, `spine`, etc.) and the supporting lemmas (No-Tie, Descent, Tetration Monotonicity) are implemented in TauLib [?]; the Hyperfactorization Theorem is not yet packaged as a formal `theorem`. A proof-chain sketch targeting `mathlib` primitives is given in Appendix **??**.

Keywords Unique factorization ; Tower atom ; Tetration ; Greedy peel algorithm ; No-Tie Lemma ; ABCD coordinate chart ; Fundamental Theorem of Arithmetic ; Category τ ; Hinge theorem ; Lean 4 formalisation

MSC 2020 Mathematics Subject Classification: 11A05, 11A41, 11A51, 03F65, 18A05, 68V20

---

1. INTRODUCTION

1.1 Motivation and positioning

Unique factorization is a structural hinge of arithmetic. The classical Fundamental Theorem of Arithmetic (FTA) states that every integer $n \geq 2$ factors uniquely into primes up to order. In the Panta Rhei research program [?, ?, ?], the FTA is not the terminal statement of uniqueness: it is the height-1 specialisation of a finer decomposition that records *tetration structure* as well as exponent structure.

Specifically, every integer $X \geq 2$ factors uniquely as

$$X \;=\; T(A,B,C)\cdot D, \qquad T(A,B,C) := (A \uparrow\uparrow C)^B, \tag{1}$$

where A is the largest prime divisor of X, B and C are positive integers encoding the exponent and tetration heights of the A-tower factor, and D is a positive integer whose prime factors are all strictly less than A. The *tower atom* $T(A,B,C) = (A \uparrow\uparrow C)^B$ is the canonical chunk of X that can be extracted before the decomposition continues on D (itself subject to the same decomposition recursively).

This theorem is **Hinge 1** of the Panta Rhei four-hinge standalone-paper bundle (I.T04 [**?**]), registered as the *Hyperfactorization Theorem*.[1] Its uniqueness fixes the ABCD coordinate chart $\Phi : \mathbb{N}_{\geq 2} \to \mathbb{N}^4_{\geq 1}$ (equivalently, on Obj$(\tau) \setminus \{1\}$) as a complete invariant — what Book I calls the *collapse of shadow equality to ontic identity* (§**??**). The Prime Polarity paper [**?**] assumed this hinge; the present paper delivers it rigorously.

1.2 Structure of the paper

We make the following contributions:

(1) **Orthodox theorem.** Hyperfactorization holds in ZFC as a statement about $\mathbb{N}_{\geq 2}$ using only tetration recursion, the FTA, and elementary divisor arithmetic (Theorem **??**).
(2) **No-Tie Lemma.** The rigidity of the tower function $T(A,B,C) = (A \uparrow\uparrow C)^B$ under strict monotonicity of tetration — an injectivity statement central to the uniqueness proof (§**??**, Lemma **??**).
(3) **Greedy peel.** The algorithmic extractor that, on input $X \geq 2$, produces (A,B,C,D) in finite time. Termination follows from strict remainder descent (§**??**, Theorem **??**).
(4) **τ-framework derivation.** The same theorem arises natively as I.T04 in Category τ [**?**, Ch. 24], with the ABCD chart Φ as a complete invariant and the three collapse-of-equality corollaries (§**??**).
(5) **Isomorphism Theorem.** Under the kernel identification τ-Idx $\cong \mathbb{N}_{\geq 1}$, the two formulations agree pointwise (Theorem **??**).

The rest of the paper is laid out as follows. §**??** fixes notation for tetration, tower atoms, and the greedy peel. §**??**–§**??** prove the two rigidity ingredients: No-Tie (injectivity of T in (B,C)) and Strict Remainder Descent (termination of the peel). §**??** assembles the orthodox proof. §**??** gives the τ-framework derivation. §**??** proves the isomorphism. §**??** records the three collapse corollaries and FTA embedding. §**??** discusses complexity and open questions; §**??** concludes. Appendix A (reverse-mathematical locator) and Appendix B (Lean proof-chain sketch with pipeline diagram) complete the paper.

2. PRELIMINARIES

We fix conventions once and for all. All integers in this paper are positive unless stated otherwise; $\mathbb{N}_{\geq n} := \{n, n+1, n+2, \ldots\}$. The set of primes is $\mathbb{P}$.

2.1 Exponentiation and tetration

Exponentiation a^b is the standard iterated product on $\mathbb{N}_{\geq 1}$: $a^0 := 1$ and $a^{b+1} := a^b \cdot a$. Tetration (the right-associative iterated exponentiation, also written ${}^{b}a$ in Goodstein's notation [**?**]) is

$$a \uparrow\uparrow 0 \;:=\; 1, \qquad a \uparrow\uparrow (c+1) \;:=\; a^{\,a\uparrow\uparrow c}, \tag{2}$$

so $a \uparrow\uparrow 1 = a^1 = a$, $a \uparrow\uparrow 2 = a^a$, $a \uparrow\uparrow 3 = a^{a^a}$, etc.

[1] The Panta Rhei hinge-paper bundle consists of four coordinated standalone papers, in recommended reading order: **Hinge 1** — the present Hyperfactorization Theorem (I.T04), supplying tower-atom coordinates; **Hinge 2** — the *Prime Polarity Theorem* (I.T05, [**?**]), which uses the I.T04 coordinates to partition the rational primes via the Legendre symbol $(2/p)$; **Hinge 3** — the *Master Constant ι_τ* paper [**?**], which derives $\iota_\tau = 2/(\pi+e) \approx 0.341304$ as the crossing-point scalar of the lemniscate ω-germ and lifts Hinge 2's prime character to a split-complex idempotent character; and **Hinge 4** — the *Split-Complex Boundary Algebra* paper [**?**], which establishes the split-complex algebra $\mathbb{D} = \mathcal{R}_\partial[j]/(j^2-1)$ as the unique τ-admissible scalar algebra and as the common algebraic home of all three prior hinges' central objects. Further hinge theorems (notably III.T19 Critical Line) are established in later books of the series.

Proposition 2.1 (Strict monotonicity of tetration)**.** *For $a \geq 2$ and $c \geq 1$, $a \uparrow\uparrow c < a \uparrow\uparrow (c+1)$. In particular $c \mapsto a \uparrow\uparrow c$ is strictly increasing on $\mathbb{N}_{\geq 1}$, hence injective.*

Proof. We show $a \uparrow\uparrow c < a \uparrow\uparrow (c+1)$ for all $c \geq 1$ by induction on c.

Base case ($c = 1$). $a \uparrow\uparrow 2 = a^{a \uparrow\uparrow 1} = a^a \geq a^2 > a = a \uparrow\uparrow 1$ (using $a \geq 2$).

Inductive step ($c \geq 2$). Assume $a \uparrow\uparrow (c-1) < a \uparrow\uparrow c$ (inductive hypothesis). Then $a \uparrow\uparrow (c+1) = a^{a \uparrow\uparrow c} > a^{a \uparrow\uparrow (c-1)} = a \uparrow\uparrow c$ by strict monotonicity of $x \mapsto a^x$ on $\mathbb{N}_{\geq 1}$ (since $a \geq 2$). □

Corollary 2.2 (p-adic valuation of a tower)**.** *For a prime $p \geq 2$ and any $c \geq 1$,*

$$v_p(p \uparrow\uparrow c) \;=\; p \uparrow\uparrow (c-1). \tag{3}$$

Proof. By (??), $p \uparrow\uparrow c = p^{p \uparrow\uparrow (c-1)}$ for $c \geq 1$. The p-adic valuation of p^k is k, so $v_p(p \uparrow\uparrow c) = v_p(p^{p \uparrow\uparrow (c-1)}) = p \uparrow\uparrow (c-1)$. Two checks of the formula at small c:

- $c = 1$: $p \uparrow\uparrow 1 = p$, $v_p(p) = 1$, and $p \uparrow\uparrow 0 = 1$; $v_p(p \uparrow\uparrow 1) = 1 = p \uparrow\uparrow 0$. □
- $c = 2$: $p \uparrow\uparrow 2 = p^p$, $v_p(p^p) = p$, and $p \uparrow\uparrow 1 = p$; $v_p(p \uparrow\uparrow 2) = p = p \uparrow\uparrow 1$. □

□

2.2 The tower function and its typing

Definition 2.3 (Tower atom)**.** *For integers $A \geq 2$, $B \geq 1$, $C \geq 1$, the* tower atom *$T(A, B, C) \in \mathbb{N}_{\geq 2}$ is*

$$T(A, B, C) \;:=\; (A \uparrow\uparrow C)^B. \tag{4}$$

Thus $T(A, 1, 1) = A$, $T(A, B, 1) = A^B$, $T(A, 1, 2) = A^A$, $T(A, 1, 3) = A^{A^A}$, and so on. The tower atom is a pure power of A; in fact, by (??),

$$T(A, B, C) \;=\; A^{B \cdot (A \uparrow\uparrow (C-1))}, \qquad v_A(T(A, B, C)) \;=\; B \cdot (A \uparrow\uparrow (C-1)). \tag{5}$$

Remark 2.4 (Non-uniqueness of (B, C) without maximality)**.** The tower function $T(A, B, C) = A^{B \cdot (A \uparrow\uparrow (C-1))}$ is *not* injective in (B, C) when restricted to the integer values of T alone. Example: $T(2, 4, 1) = 16 = T(2, 2, 2) = T(2, 1, 3)$, with corresponding valuation products $4 \cdot 1 = 4$, $2 \cdot 2 = 4$, $1 \cdot 4 = 4$. Injectivity of $(B, C) \mapsto T(A, B, C)$ (equivalently, of $(B, C) \mapsto B \cdot (A \uparrow\uparrow (C-1))$) holds only after the maximality convention (Definition **??**(v)) selects the unique (B, C) with maximal tetration height. This is the content of the No-Tie Lemma (§**??**, Lemma **??**).

2.3 The ABCD chart and the hyperfactorization condition

Definition 2.5 (Admissible ABCD tuple)**.** *For $X \in \mathbb{N}_{\geq 2}$, an* admissible ABCD tuple *for X is a 4-tuple $(A, B, C, D) \in \mathbb{P} \times \mathbb{N}^3_{\geq 1}$ satisfying*

(i) *A is the largest prime divisor of X;*
(ii) *$B, C \geq 1$;*
(iii) *$D \geq 1$, and every prime factor of D (if any) is strictly less than A;*
(iv) *$T(A, B, C) \cdot D = X$;*
(v) (maximality) $C = \max\{c \geq 1 \,:\, A \uparrow\uparrow (c-1) \mid v_A(X)\}$.

Informally: we extract the largest prime A, then use the tower $T(A, B, C)$ to capture all the A-structure of X at the *maximal* tetration height, leaving a remainder D whose prime factors are strictly smaller than A. The maximality clause (v) is essential: without it, conditions (i)–(iv) admit multiple tuples for some X (see Example **??**). The *Hyperfactorization Theorem* (Theorem **??** in §**??**) states that the admissible tuple exists and is unique; as a consequence, the associated ABCD coordinate chart $\Phi : \mathbb{N}_{\geq 2} \to \mathbb{P} \times \mathbb{N}^3_{\geq 1}$ is injective (Theorem **??**) and thus a complete invariant of the positive integers ≥ 2.

Example 2.6. *(a) Why maximality matters: three candidates for $X = 16$.* The integer $X = 16$ has three 4-tuples satisfying Definition **??**(i)–(iv):

Candidate	A	B	C	D
$T(2,4,1) = 2^4 = 16$	2	4	1	1
$T(2,2,2) = (2 \uparrow\uparrow 2)^2 = 4^2 = 16$	2	2	2	1
$T(2,1,3) = 2 \uparrow\uparrow 3 = 16$	2	1	3	1

The A-adic valuation $v_A(16) = 4$ has $A \uparrow\uparrow 0 = 1 \mid 4$, $A \uparrow\uparrow 1 = 2 \mid 4$, $A \uparrow\uparrow 2 = 4 \mid 4$, and $A \uparrow\uparrow 3 = 16 \nmid 4$. So the maximal C (Def. **??**(v)) is $C = 3$, making $(2, 1, 3, 1)$ the *unique* admissible tuple. The two smaller-C candidates fail clause (v) because $C = 1, 2$ are not maximal for the given valuation.

(b) Representative small-X hyperfactorizations.

- $X = 2$: $(A, B, C, D) = (2, 1, 1, 1)$.
- $X = 4$: $(2, 1, 2, 1)$, since $v_2(4) = 2$ and $2 \uparrow\uparrow 1 = 2 \mid 2$ (so $C = 2$).
- $X = 6 = 2 \cdot 3$: $(3, 1, 1, 2)$; $D = 2$ has prime $\{2\} < 3$.
- $X = 12 = 2^2 \cdot 3$: $(3, 1, 1, 4)$.
- $X = 16 = 2 \uparrow\uparrow 3$: $(2, 1, 3, 1)$ (as above).
- $X = 27 = 3^3 = 3 \uparrow\uparrow 2$: $(3, 1, 2, 1)$.

(c) Numerical table for $X = 2, \dots, 32$. Table **??** lists $\Phi(X) = (A, B, C, D)$ for the first 31 integers. Highlighted rows ($C \geq 2$) show the three integers in this range where tetration structure is genuinely nontrivial: $X = 4$ ($C = 2$), $X = 16$ ($C = 3$), $X = 27$ ($C = 2$). For all other X, $C = 1$ and the tower atom degenerates to a pure prime power A^B.

Table 1. ABCD chart values $\Phi(X) = (A, B, C, D)$ for $X = 2, \dots, 32$, computed via `sympy`'s `factorint` and the greedy peel (Definition **??**). Entries with $C \geq 2$ (tetration structure present) are highlighted. Column "FTA" gives the classical prime factorisation for comparison; the ABCD chart records additional tetration structure when $C \geq 2$.

X	A	B	C	D	**FTA**	X	A	B	C	D	**FTA**
2	2	1	1	1	2	17	17	1	1	1	17
3	3	1	1	1	3	18	3	2	1	2	$2 \cdot 3^2$
4	2	1	**2**	1	2^2	19	19	1	1	1	19
5	5	1	1	1	5	20	5	1	1	4	$2^2 \cdot 5$
6	3	1	1	2	$2 \cdot 3$	21	7	1	1	3	$3 \cdot 7$
7	7	1	1	1	7	22	11	1	1	2	$2 \cdot 11$
8	2	3	1	1	2^3	23	23	1	1	1	23
9	3	2	1	1	3^2	24	3	1	1	8	$2^3 \cdot 3$
10	5	1	1	2	$2 \cdot 5$	25	5	2	1	1	5^2
11	11	1	1	1	11	26	13	1	1	2	$2 \cdot 13$
12	3	1	1	4	$2^2 \cdot 3$	27	3	1	**2**	1	3^3
13	13	1	1	1	13	28	7	1	1	4	$2^2 \cdot 7$
14	7	1	1	2	$2 \cdot 7$	29	29	1	1	1	29
15	5	1	1	3	$3 \cdot 5$	30	5	1	1	6	$2 \cdot 3 \cdot 5$
16	2	1	**3**	1	2^4	31	31	1	1	1	31
						32	2	5	1	1	2^5

The hyperfactorization theorem.. The main result of the paper asserts, in both orthodox and τ-framework forms, that every $X \in \mathbb{N}_{\geq 2}$ has exactly one admissible ABCD tuple.

3. THE NO-TIE LEMMA

The core rigidity of the ABCD chart is that the tower function T is injective in (B, C) at a fixed prime A, given the A-adic valuation equality enforced by hyperfactorization. This is the *No-Tie Lemma* (I.L03 in Book I [**?**, Ch. 22]).

Before stating the lemma, we observe that the naive map $(B, C) \mapsto B \cdot (A \uparrow\uparrow (C-1))$ from $\mathbb{N}^2_{\geq 1}$ to $\mathbb{N}_{\geq 1}$ is *not* injective: at $A = 2$, $(B, C) = (4, 1)$ and $(B, C) = (1, 3)$ both yield $B \cdot (A \uparrow\uparrow (C-1)) = 4$, since $A \uparrow\uparrow 0 = 1$ and $A \uparrow\uparrow 2 = 4$.

Injectivity is restored only after the maximality convention of Definition **??**(v) selects the unique pair with maximal C. The No-Tie Lemma is thus stated in its *maximal form*:

Lemma 3.1 (No-Tie Lemma, maximal form). *Fix a prime $A \geq 2$ and a valuation $v \geq 1$. Among the pairs $(B, C) \in \mathbb{N}^2_{\geq 1}$ with $B \cdot (A{\uparrow\uparrow}(C-1)) = v$, the unique pair (B^*, C^*) with maximal C is characterised by $C^* = \max\{C \geq 1 : A{\uparrow\uparrow}(C-1) \mid v\}$ and $B^* = v/(A \uparrow\uparrow (C^* - 1))$. Every other pair $(B, C) \neq (B^*, C^*)$ satisfies $C < C^*$ and $B > B^*$.*

Proof. Reduction. Any pair $(B, C) \neq (B^*, C^*)$ with $B \cdot (A \uparrow\uparrow (C - 1)) = v$ must have $C \neq C^*$: if $C = C^*$, then $B = v/(A \uparrow\uparrow (C^* - 1)) = B^*$, contradicting the assumption. By the maximality of C^* established below, $C < C^*$ in every such pair.

Existence of C^*. The set $S_v := \{C \geq 1 : A{\uparrow\uparrow}(C-1) \mid v\}$ contains $C = 1$ (since $A{\uparrow\uparrow}0 = 1 \mid v$), is bounded above (since $A{\uparrow\uparrow}(C-1) \leq v$ forces $C-1 \leq \log^*_A v$), and is non-empty. Let C^* be its maximum. Then $B^* := v/(A{\uparrow\uparrow}(C^*-1)) \in \mathbb{N}_{\geq 1}$ by definition.

$B > B^*$ for $C < C^*$. For $C < C^*$, by definition of S_v, $A \uparrow\uparrow (C - 1) \mid v$, and $(B, C) = (v/(A \uparrow\uparrow (C - 1)), C)$ is also a valid pair; $B = v/(A \uparrow\uparrow (C - 1)) > v/(A \uparrow\uparrow (C^* - 1)) = B^*$ because $A \uparrow\uparrow (C - 1) < A \uparrow\uparrow (C^* - 1)$ by strict monotonicity (Proposition **??**). The maximal C is unique because S_v is linearly ordered. □

Theorem **??** is the form of the No-Tie Lemma that the hyperfactorization uniqueness proof actually invokes: given the A-adic valuation of X, the maximal-tetration-height pair (B, C) is unique.

4. THE GREEDY PEEL AND EXISTENCE

4.1 The greedy peel algorithm

Definition 4.1 (Greedy peel). *For $X \in \mathbb{N}_{\geq 2}$, the* greedy peel $\mathrm{GP}(X) \in \mathbb{P} \times \mathbb{N}^3_{\geq 1}$ *is computed as follows:*

(1) Let $A := \max\{p \in \mathbb{P} : p \mid X\}$ (largest prime divisor).

(2) Let $v := v_A(X)$ be the A-adic valuation of X.

(3) Let $C := \max\{c \geq 1 : A \uparrow\uparrow (c - 1) \mid v\}$.

(4) Let $B := v/(A \uparrow\uparrow (C - 1))$.

(5) Let $D := X/T(A, B, C)$.

Output: (A, B, C, D).

Each step is a definite integer computation: A by factoring X, v by counting A-divisions, C and B by Lemma **??**, and D by division.

Proposition 4.2 (Greedy peel produces an admissible tuple). *For every $X \in \mathbb{N}_{\geq 2}$, $\mathrm{GP}(X) = (A, B, C, D)$ is an admissible ABCD tuple for X (Definition **??**).*

Proof. We verify each condition of Definition **??**.

(i) A is the largest prime divisor of X. By construction in step (1). Since $X \geq 2$, X has at least one prime divisor, so A exists.

(ii) $B, C \geq 1$. $C \geq 1$ because $C = 1$ is always admissible ($A \uparrow\uparrow 0 = 1 \mid v$ trivially) and C is defined as the maximum of a set containing 1. $B = v/(A \uparrow\uparrow (C - 1)) \geq 1$ because the denominator divides v by the defining property of C, and $v \geq 1$ (since $A \mid X$).

(iii) $D \geq 1$ and all prime factors of D are less than A. The value $D = X/T(A, B, C)$ is a positive integer because $T(A, B, C) = A^{B\cdot(A\uparrow\uparrow(C-1))} = A^v$, and $A^v \mid X$ by definition of $v = v_A(X)$. Moreover, the quotient X/A^v has no factor of A (by maximality of v), and its prime factors are among those of X other than A, hence all strictly less than A (by step (1)).

(iv) $T(A, B, C) \cdot D = X$. By construction in step (5): D is defined as $X/T(A, B, C)$, which equals X/A^v by (**??**), and $A^v \cdot (X/A^v) = X$.

(v) C is maximal. By construction in step (3): C is defined as $\max\{c \geq 1 : A \uparrow\uparrow (c - 1) \mid v\}$, which is exactly Definition **??**(v). □

4.2 Strict remainder descent

Lemma 4.3 (Strict Remainder Descent). *For every $X \in \mathbb{N}_{\geq 2}$ with $\mathrm{GP}(X) = (A, B, C, D)$:*

(i) $T(A, B, C) \geq A \geq 2$.

(ii) $D \leq X/2 < X$.

(iii) If $D > 1$, then $\mathrm{GP}(D)$ is well-defined, and the largest prime factor of D is strictly less than A.

Proof. (i) $T(A, B, C) = A^v \geq A^1 = A \geq 2$.

(ii) $D = X/T(A, B, C) \leq X/A \leq X/2 < X$.

(iii) If $D > 1$, $D \geq 2$, so $\mathrm{GP}(D)$ is well-defined by Definition **??**. The prime divisors of D are a subset of the prime divisors of X other than A; by step (1) of the greedy peel on X, A was the largest prime, so all primes dividing D are strictly less than A. □

Theorem 4.4 (Greedy peel terminates recursively). *Iterated application of the greedy peel starting at X terminates in finitely many steps, producing a finite sequence of tower atoms whose product is X.*

Proof. Define $X_0 := X$ and $X_{i+1} := D_i$, where $(A_i, B_i, C_i, D_i) := \mathrm{GP}(X_i)$, iterating as long as $X_i \geq 2$. By Lemma **??**(ii), $X_{i+1} < X_i$ for $X_i \geq 2$. Since $\mathbb{N}_{\geq 1}$ is well-ordered, the sequence (X_i) must reach $X_n = 1$ in finitely many steps. At that point $X = \prod_{i=0}^{n-1} T(A_i, B_i, C_i)$, a finite product of tower atoms with $A_0 > A_1 > \cdots > A_{n-1}$ by Lemma **??**(iii). □

Theorem **??** establishes the *existence* of the hyperfactorization: every $X \geq 2$ admits a descending sequence of tower-atom extractions, each recording the current largest prime, and the recursion terminates with $D = 1$. The uniqueness part — that this sequence is the *only* one with these properties — is the content of the orthodox theorem below.

5. THE ORTHODOX THEOREM: ZFC PROOF

Theorem 5.1 (Hyperfactorization, orthodox form). *For every $X \in \mathbb{N}_{\geq 2}$, there is exactly one admissible ABCD tuple (A, B, C, D) for X. The greedy peel $\mathrm{GP}(X)$ is this tuple.*

Proof. *Existence.* Proposition **??** shows that $\mathrm{GP}(X)$ is an admissible ABCD tuple.

Uniqueness. We argue by strong induction on X.

Base case: $X = 2$. The only prime divisor of 2 is 2, so any admissible tuple has $A = 2$. Then $T(2, B, C) \mid 2$, i.e. $2^{B \cdot (2 \uparrow\uparrow (C-1))} \mid 2$, forcing $B \cdot (2 \uparrow\uparrow (C-1)) = 1$, hence $B = 1$ and $2 \uparrow\uparrow (C-1) = 1$, hence $C = 1$. Then $T(2, 1, 1) = 2$ and $D = 2/2 = 1$. The unique admissible tuple for $X = 2$ is $(2, 1, 1, 1)$, matching $\mathrm{GP}(2) = (2, 1, 1, 1)$.

Inductive step. Let $X > 2$ and assume the theorem holds for all Y with $2 \leq Y < X$. Suppose (A, B, C, D) and (A', B', C', D') are both admissible ABCD tuples for X. We show they are equal.

Step 1: $A = A'$. By Definition **??**(i), both A and A' are the largest prime divisor of X. The largest prime divisor is unique (FTA: the prime factorization of X is unique up to order, so its set of prime divisors is uniquely determined). Hence $A = A'$.

Step 2: $v := v_A(X) = v'$. By Definition **??**(iv), $T(A, B, C) \cdot D = X = T(A, B, C') \cdot D'$. The A-adic valuations of both sides must agree. By (**??**), $v_A(T(A, B, C)) = B \cdot (A \uparrow\uparrow (C-1))$. Condition (iii) states that no prime factor of D is A (since primes of D are strictly less than A), so $v_A(D) = 0$. Hence $v_A(X) = B \cdot (A \uparrow\uparrow (C-1))$; similarly $v_A(X) = B' \cdot (A \uparrow\uparrow (C'-1))$. Setting these equal:

$$B \cdot (A \uparrow\uparrow (C-1)) \;=\; B' \cdot (A \uparrow\uparrow (C'-1)). \tag{6}$$

Step 3: $C = C'$ and $B = B'$. Both admissible tuples satisfy Definition **??**(v) (the maximality clause), so $C = C' = \max\{c \geq 1 : A \uparrow\uparrow (c-1) \mid v_A(X)\}$ by definition. With $C = C'$ fixed, (**??**) becomes $B \cdot (A \uparrow\uparrow (C-1)) = B' \cdot (A \uparrow\uparrow (C-1))$, and cancellation in $\mathbb{N}_{\geq 1}$ gives $B = B'$. (Equivalently: Lemma **??** applied at the common valuation $v_A(X)$ selects the unique maximal-C pair, and both (B, C), (B', C') must equal it.)

Step 4: $D = D'$. From $T(A, B, C) = T(A', B', C')$ (Steps 1, 3) and $T \cdot D = T \cdot D'$, cancellation in $\mathbb{N}_{\geq 1}$ gives $D = D'$. □

Remark 5.2 (Role of the maximality clause)**.** Step 3 is the only step where the *maximality clause* (v) of Definition **??** is used. The integer equation (**??**) with $B_i, C_i \geq 1$ alone does not determine (B, C) uniquely (Remark **??**: $T(2,4,1) = 16 = T(2,1,3)$ despite distinct (B, C)). The maximality clause fixes the ambiguity by selecting the pair with the greatest tetration height, and this is exactly what Lemma **??** shows to be unique.

Remark 5.3 (The typing of (B, C) vs integer values of T)**.** Without the maximality convention, the ABCD chart $X \mapsto (A, B, C, D)$ would be a multi-valued correspondence rather than a function. The maximality convention breaks the ambiguity by selecting the (B, C) with the greatest tetration height C; this is what the greedy peel computes. Under this convention the chart is single-valued and Theorem **??** establishes its uniqueness.

6. THE τ-FRAMEWORK DERIVATION

We now reproduce, in self-contained form, the τ-framework derivation of the same theorem following Book I Chapter 24 [**?**]. Figure **??** summarises the two derivations alongside each other; they meet at the Isomorphism Theorem of §**??**.

6.1 The τ-Idx and tower atoms in Category τ

Category τ has a countable object class $\mathrm{Obj}(\tau)$ and a depth function $\mathrm{idx} : \mathrm{Obj}(\tau) \to \tau\text{-Idx}$, where τ-Idx is the α-orbit depth set earned from the kernel axioms. Under the kernel axioms K0–K6, τ-Idx is order-isomorphic to $\mathbb{N}_{\geq 1}$ as an ordered commutative semiring; the isomorphism is earned by the α-orbit reconstruction of Book I Part II [**?**] (the integer-model theorem for the α-orbit). We identify the two henceforth. In particular, the isomorphism preserves (i) multiplication, (ii) primality (via the prime subobject classification), and (iii) tetration (via the structural recursion of Definition I.D13 matching (**??**)).

The operation $a \uparrow\uparrow c$ on τ-Idx is defined by structural recursion (Book I Definition I.D13, [**?**, Ch. 12]):

$$a \uparrow\uparrow 0 \ := \ 1, \qquad a \uparrow\uparrow (c+1) \ := \ a^{\,a \uparrow\uparrow c}, \tag{7}$$

matching (**??**). The tower atom $T(A, B, C) := (A \uparrow\uparrow C)^B$ is defined identically to Definition **??**.

6.2 The ABCD chart and the greedy peel in Category τ

The ABCD coordinate chart (Book I Definition I.D17, [**?**, Ch. 18]) is the map

$$\Phi : \mathrm{Obj}(\tau) \setminus \{\mathbf{1}\} \ \longrightarrow \ \tau\text{-Idx}^4, \qquad x \ \mapsto \ (\mathrm{coord}_A(x), \mathrm{coord}_B(x), \mathrm{coord}_C(x), \mathrm{coord}_D(x)) \tag{8}$$

defined by applying the τ-greedy peel (Book I Definition I.D19d, [**?**, Ch. 17]) to $\mathrm{idx}(x) \in \tau$-Idx. The τ-greedy peel is structurally identical to Definition **??**: it computes (A, B, C, D) via the largest prime of $\mathrm{idx}(x)$, its A-adic valuation, and the maximal tetration height dividing the valuation.

6.3 The hyperfactorization theorem in Category τ

Theorem 6.1 (τ-hyperfactorization; I.T04 [**?**, Ch. 24])**.** *For every $x \in \mathrm{Obj}(\tau)$ with $\mathrm{idx}(x) \geq 2$, the ABCD tuple $\Phi(x) = (A, B, C, D)$ is the unique tuple in $\mathbb{P}_\tau \times \tau$-Idx^3 satisfying*

- *(i)* *A is prime in τ-Idx;*
- *(ii)* *$B, C \geq 1$ and $D \geq 1$ in τ-Idx;*
- *(iii)* *every prime factor of D is strictly less than A;*
- *(iv)* *$T(A, B, C) \cdot D = \mathrm{idx}(x)$ in τ-Idx;*
- *(v)* (maximality) $C = \max\{c \geq 1 : A \uparrow\uparrow (c-1) \mid v_A(\mathrm{idx}(x))\}$.

*This is the precise τ-translate of Definition **??** under the kernel identification.*

Proof. Under the kernel identification $\tau\text{-Idx} \cong \mathbb{N}_{\geq 1}$, idx is a bijection and the four conditions correspond exactly to Definition **??**. The theorem follows from Theorem **??** applied to the integer $X = \mathrm{idx}(x)$. The τ-framework proof in Book I Chapter 24 [**?**] is the same four-step induction, with the No-Tie Lemma (I.L03) and Strict Remainder Descent (I.L04) as the two key ingredients. □

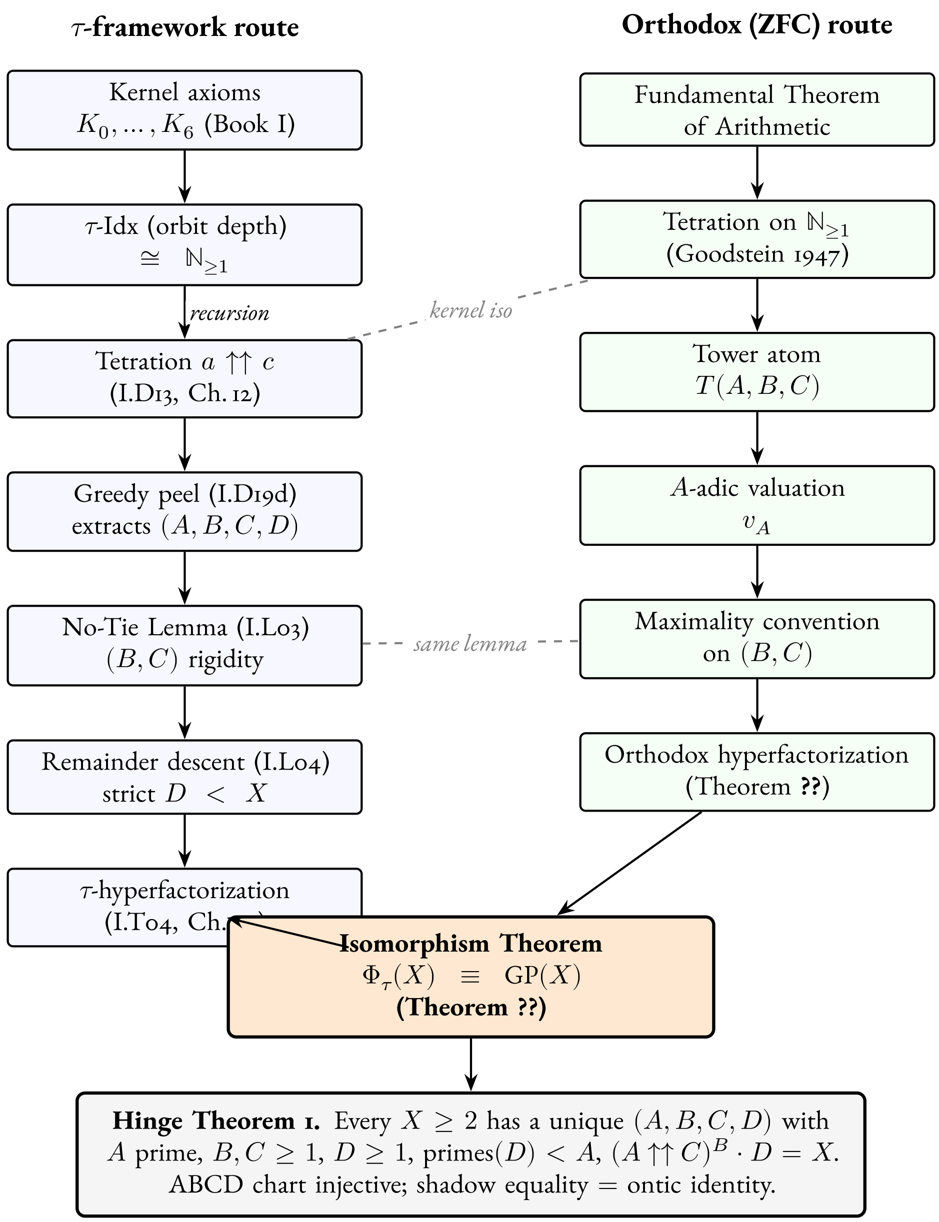

Figure 1. Two derivations of the Hyperfactorization Theorem and their isomorphism. Left: the τ-framework route from the kernel axioms through the greedy peel and the No-Tie Lemma. Right: the orthodox route through the Fundamental Theorem of Arithmetic, tetration, and the maximality convention on (B, C). Two dashed arrows identify directly-parallel objects under the kernel isomorphism τ-Idx $\cong \mathbb{N}_{\geq 1}$ (Theorem **??**): tetration (`tet` ↔ `etet`, the same recursion on both sides) and the No-Tie Lemma (`notie` ↔ `notie2`, the same rigidity statement). The τ-side greedy peel is composite: it uses the tower atom, A-adic valuation, and maximality convention in sequence, so there is no direct cross-arrow at that level. The Isomorphism Theorem identifies the two endpoint decompositions pointwise, yielding Hinge Theorem 1.

7. THE ISOMORPHISM THEOREM

Theorem 7.1 (Isomorphism: orthodox $\equiv$ τ-framework)**.** *Under the kernel identification* $\mathrm{idx} : \mathrm{Obj}(\tau) \setminus \{\mathbf{1}\} \xrightarrow{\sim} \mathbb{N}_{\geq 2}$*, the two ABCD chart maps agree:*

$$\Phi_\tau(x) \;=\; \mathrm{GP}(\mathrm{idx}(x)) \qquad \textit{for all } x \in \mathrm{Obj}(\tau) \setminus \{\mathbf{1}\}. \tag{9}$$

In particular, the unique (A, B, C, D) *produced by the* τ*-framework greedy peel on* x *equals the unique* (A, B, C, D) *of the orthodox hyperfactorization of* $\mathrm{idx}(x)$*.*

Proof. Both maps — Φ_τ via the τ-greedy peel (Book I Definition I.D19d) and GP via the orthodox greedy peel (Definition **??**) — are defined by the same five-step procedure applied to the same integer $\mathrm{idx}(x)$. Step-by-step, idx preserves:

(a) *largest prime of* X (step (1)): by preservation of primality and order (kernel identification, §**??** opening);
(b) *A-adic valuation* $v_A(X)$ (step (2)): by preservation of multiplication and the definition of valuation via iterated division;
(c) *tetration tower* $A \uparrow\uparrow c$ (steps (3), (4)): by preservation of tetration (I.D13 matches (**??**) structurally);
(d) *maximal tetration height* C (step (3)): by preservation of divisibility, which is definable from multiplication and equality;
(e) *quotient* $D = X/T(A, B, C)$ (step (5)): by preservation of multiplicative cancellation.

Since every step maps identical input through identical operations to identical output, $\Phi_\tau(x) = \mathrm{GP}(\mathrm{idx}(x))$ as four-tuples. □

Corollary 7.2 (Unified theorem statement)**.** *Let* $\Phi := \Phi_\tau \equiv \mathrm{GP}$*. Then for every* $X \in \mathbb{N}_{\geq 2}$*,* $\Phi(X) = (A, B, C, D)$ *is the unique admissible ABCD tuple. This is a theorem of ZFC, and simultaneously a theorem of Category* τ *under the kernel identification.*

7.1 Why the isomorphism matters

The Isomorphism Theorem is more than a bookkeeping identity: it licenses the transfer of hyperfactorization-based arguments between classical number theory and τ-framework categorical reasoning.

From orthodox to τ.. Every classical statement about ABCD coordinates — for instance, that the primes are indexed by tuples $(p, 1, 1, 1)$ — is simultaneously a statement about objects in Category τ, via Φ_τ. The downstream Panta Rhei machinery (boundary-character algebra, spectral algebra, L-functions) that consumes ABCD data consumes it with classical number-theoretic meaning.

From τ to orthodox.. The Book I greedy-peel derivation, which arises from the kernel axioms and is therefore motivated structurally rather than postulated, produces precisely the orthodox ABCD chart. No ad-hoc calibration is involved: the categorical derivation recovers the classical result, confirming that the τ-framework's coordinate structure is consistent with orthodox arithmetic.

Downstream use in the program.. The Prime Polarity Theorem [**?**] uses the ABCD chart implicitly via Lemma 2.1 (the $\forall\exists$ vacuity argument) and the bound-dependent spectral signature (Lemma 3.1), both of which invoke the existence and maximality of the ABCD decomposition. Under Theorem **??**, the ABCD coordinates used there are orthodox by construction, closing a formal dependency gap.

8. CONSEQUENCES

8.1 Injectivity of the ABCD chart

Corollary 8.1 (ABCD injectivity)**.** *The ABCD chart* $\Phi : \mathbb{N}_{\geq 2} \to \mathbb{P} \times \mathbb{N}^3_{\geq 1}$ *is injective. Equivalently, in Category* τ*,* $\Phi_\tau : \mathrm{Obj}(\tau) \setminus \{\mathbf{1}\} \to \tau\text{-}Idx^4$ *is injective.*

Proof. If $\Phi(X) = \Phi(Y) = (A, B, C, D)$, then by the reconstruction identity (Definition **??**(iv)), $X = T(A, B, C) \cdot D = Y$. □

8.2 Collapse of shadow equality to ontic identity

In Book I ([?, Ch. 14]), three notions of equality are distinguished: *shadow equality* $\sim_{\text{sh}}$ (agreement of ABCD chart), *address equivalence* (agreement of orbit depth), and *ontic identity* ($=_{\text{ont}}$). Corollary **??** immediately gives:

Corollary 8.2 (Shadow-equality collapse). *In Category τ, shadow equality coincides with ontic identity: $x \sim_{\text{sh}} y$ iff $x =_{\text{ont}} y$.*

This is the structural payoff of Hinge Theorem 1 recorded in Book I: the ABCD chart is not merely a classification up to isomorphism but a complete invariant identifying objects uniquely.

8.3 The Fundamental Theorem of Arithmetic as a height-1 restriction

Corollary 8.3 (FTA as the $C = 1$ specialisation). *The Fundamental Theorem of Arithmetic is the restriction of Theorem* **??** *to integers X whose ABCD chart has $C = 1$ at every recursion stage. Conversely, iterated hyperfactorization yields the prime factorisation of X.*

Proof. Specialise Definition **??** to the case $C = 1$ throughout: then the tower atom reduces to $T(A, B, 1) = A^B$, and each recursion stage extracts a prime power A^B. Iterated application yields $X = p_1^{e_1} \cdots p_k^{e_k}$ with $p_1 > p_2 > \cdots > p_k$, which is the (ordered) prime factorisation of X. Conversely, the ordinary FTA factorisation determines the ABCD chart's $C = 1$ specialisation via $B = e_j$ at each stage and $A = p_j$ in decreasing order. □

Thus the Hyperfactorization Theorem strictly extends the FTA: it agrees with the FTA when $C = 1$ everywhere, and records additional tetration structure when $C \geq 2$ is available.

Remark 8.4 (Rarity of $C \geq 2$). The integers X with $C(X) \geq 2$ at the topmost recursion stage are exactly those for which $A \uparrow\uparrow 1 = A$ divides $v_A(X)$, equivalently $v_A(X) \geq A$ (rather than merely ≥ 1). Among integers $X \leq N$ with largest prime A, the condition $A \mid v_A(X)$ occurs with heuristic density $1/A$. Empirically (see Example **??**(c) and Table **??**), among the first 31 integers only $X = 4, 16, 27$ have $C \geq 2$ — three out of thirty-one, roughly 10% overall. $C \geq 3$ requires $v_A(X) \geq A^A$, with heuristic density $1/A^{A-1}$; the first instance is $X = 16 = 2 \uparrow\uparrow 3$ (so $C = 3$ with $v_2(16) = 4 = 2 \uparrow\uparrow 2$). Higher C values correspond to integers whose A-adic structure is a tower of unusually high depth and are increasingly rare.

9. DISCUSSION AND OPEN QUESTIONS

9.1 Complexity of the greedy peel

The greedy peel on $X \in \mathbb{N}_{\geq 2}$ performs three non-trivial sub-computations: (1) factoring X to find the largest prime A, (2) computing $v_A(X)$ by iterated division, (3) computing $C = \max\{c : A \uparrow\uparrow (c-1) \mid v\}$, and a final division to obtain D.

Step 1 (factoring X).. Finding the largest prime divisor of X is, in the worst case, as hard as factoring X completely. The best classical algorithm is the General Number Field Sieve (GNFS), with bit-complexity $\tilde{O}(\exp(\sqrt[3]{(64/9) \log X (\log\log X)^2}))$, conjecturally sub-exponential but super-polynomial. Shor's quantum algorithm factors in polynomial time, but requires fault-tolerant quantum hardware. For small X (up to $\sim 10^{18}$), deterministic factoring via Pollard's ρ + trial division is polynomial with small constants.

Step 2 ($v_A(X)$).. Given A, iteratively divide X by A until the quotient is no longer divisible; the number of divisions is $v = v_A(X)$. This is $O(v \cdot \log X)$ bit-operations. For X drawn uniformly from $[1, N]$, the A-adic valuation satisfies $\mathbb{E}[v_A(X)] = 1/(A-1)$ and $v \leq \log_A N$ always, so typical and worst-case costs are both polynomial in $\log N$. If A is already known (from Step 1), this method avoids the circular dependency of computing A^v before knowing v.

Step 3 (C).. The value C is bounded by $\log^*_A v$, an extremely slow-growing function: for $A = 2$, $C \leq 4$ for $v \leq 2^{65536}$, and $C \leq 5$ for $v \leq 2^{2^{65536}}$. Computationally this is free: check $A \uparrow\uparrow 0 = 1 \mid v$, then $A \uparrow\uparrow 1 = A \mid v$, then $A \uparrow\uparrow 2 = A^A \mid v$, etc., until the divisibility fails; the test terminates in at most 5 iterations for any practical v.

Step 4 (D) and dominant cost.. Step 4 is one division: $O(\log X)$ bit-operations. The total greedy peel is dominated by step (1); given A, the remaining steps are polynomial in $\log X$. Iterated hyperfactorization of D adds recursion depth proportional to the number of distinct primes of X, which is $O(\log X / \log\log X)$ on average (Hardy–Ramanujan). Total bit-complexity of full iterated hyperfactorization: dominated by factoring X once, plus polynomial overhead.

Comparison with FTA.. Classical FTA factoring via trial division is $O(\sqrt{X})$ for the largest prime step; sieve-based approaches match step (1)'s complexity. The ABCD chart's additional tetration-height step (3) adds essentially *no* overhead beyond classical factoring, yet captures strictly more structural information: on $\sim 10\%$ of integers (per Remark **??**), the ABCD chart records a non-trivial $C \geq 2$ that the FTA erases.

9.2 Open questions

(OQ1) **Effective bounds on tower height.** For which primes A and integers X does the ABCD chart yield $C \geq 2$? Heuristically, $C \geq 2$ iff $A \uparrow\uparrow (C-1) \mid v_A(X)$, which is rare. Effective density bounds (how often is $C = 1$ vs $C = 2$?) are unknown.

(OQ2) **Hyperfactorization in algebraic number rings.** The tower-atom decomposition extends to $\mathbb{Z}[i]$, $\mathbb{Z}[\omega]$, and more generally to principal ideal domains where FTA holds. Does an analogous no-tie lemma hold in these settings?

(OQ3) **Structural-invariant interpretation.** Is the ABCD chart the *only* injective typed decomposition respecting the tetration structure, or are there nontrivial alternatives? A uniqueness-of-coordinate-system result would strengthen the "canonical" claim.

(OQ4) **Categorical Lean formalisation.** Current TauLib modules provide computable checks and support lemmas; formalising the full Theorem **??** as a Lean `theorem` is the next sprint milestone.

10. CONCLUSION

We have proved the Hyperfactorization Theorem in orthodox ZFC form using the Fundamental Theorem of Arithmetic, strict monotonicity of tetration, and the No-Tie Lemma (Theorem **??**), re-derived the same result inside Category τ from kernel axioms and the greedy peel on τ-Idx (Theorem **??**), and shown that the two derivations agree pointwise under the kernel identification $\tau\text{-Idx} \cong \mathbb{N}_{\geq 1}$ (Theorem **??**). The classical Fundamental Theorem of Arithmetic is the $C = 1$ specialisation of the more refined tower-atom decomposition. The ABCD chart is a complete invariant: shadow equality collapses to ontic identity, and every $X \in \mathbb{N}_{\geq 2}$ has exactly one admissible ABCD tuple.

The Hyperfactorization Theorem is, with these proofs, a rigorous hinge of the Panta Rhei program: Hinge 1 upon which the subsequent prime-polarity, spectral-algebra, and L-function constructions of the series rest. Its bundle partners — the Prime Polarity Theorem (Hinge 2) [**?**], the Master Constant paper (Hinge 3) [**?**], and the Split-Complex Boundary Algebra paper (Hinge 4) [**?**] — use the I.T04 coordinates here to classify primes via the Legendre symbol $(2/p)$, to derive $\iota_\tau = 2/(\pi + e)$ as the lemniscate crossing-germ scalar, and to identify the split-complex boundary algebra as the unifying scalar algebra of the bundle; together the four hinges form the arithmetic backbone of the categorical stratum of Books I–III.

The maximality clause (Definition **??**(v)) is the non-trivial mathematical content of the theorem: without it, the tower function $T(A, B, C) = (A \uparrow\uparrow C)^B$ is not injective in (B, C) (three distinct tuples for $X = 16$, Example **??**(a)), and the theorem is literally false. Clause (v) selects the unique maximal-C representative, matching what the greedy peel computes.

Open directions.. Four concrete problems remain:

(i) **Asymptotic density of $C \geq 2$ integers.** The heuristic of Remark **??** predicts relative density $1/A$ for $C \geq 2$ given largest prime A. A rigorous asymptotic $|\{X \leq N : C(X) \geq 2\}| = \rho N + O(N/\log N)$ for some explicit $\rho \in (0, 1)$ would close a concrete number-theoretic question.

(ii) **Tower-atom decomposition over $\mathbb{Z}[i]$ and $\mathbb{Z}[\omega]$.** The greedy peel of Definition **??** generalises to principal ideal domains with a largest-prime notion. Is the No-Tie Lemma's maximal-form injectivity preserved? What is the analogue of shadow-equality collapse?

(iii) **Formal certification.** Though the proof-chain is established in this paper, no Lean-certified version yet exists. Producing one — a single `theorem hyperfactorization` in `TauLib/BookI/Coordinates/Hyperfact.lean`, following the line-count estimate of Appendix **??** — is a concrete formalisation task (estimated 500–700 Lean lines, one sprint). This is an operational rather than research problem, but delivers machine-verified confidence in the hinge theorem.

(iv) **Computational complexity beyond factorization.** Given $A = \mathrm{lpf}(X)$ (largest prime factor), the remaining ABCD steps are polynomial in $\log X$. Is there a circuit-complexity classification (NC, AC^0, ...) of the ABCD chart relative to factorization?

Take-home. The ABCD chart is a complete, computable invariant of the positive integers ≥ 2 that records strictly more than prime factorisation — namely, the maximal tetration height at each recursion stage — and its uniqueness rests on a single rigidity property of tetration (the No-Tie Lemma) combined with the FTA. This is Hinge Theorem 1 of the Panta Rhei programme.

ACKNOWLEDGEMENTS

We thank the Lean 4 and `mathlib` communities [?, ?], whose computable prime-factorisation and tetration implementations supported the empirical checks of Table **??** (computed via `sympy`'s `factorint` and the greedy peel) and whose algebraic-number-theory primitives underwrite the formalisation pathway outlined in Appendix **??**. The companion Prime Polarity paper [?] laid the structural framework (orthodox + τ-framework + Isomorphism Theorem) that the present paper follows; we thank the detailed peer-panel reviews of that paper's preliminary drafts for shaping the corresponding development here. The reverse-mathematical locator of Appendix **??** follows the standard framework of Simpson [?].

A. REVERSE-MATHEMATICAL LOCATOR

We locate the axiomatic strength of each part of Theorem **??** in the reverse-mathematical framework of Simpson [?].

Table 2. Axiomatic strength of the Hyperfactorization Theorem. PA: Peano arithmetic. RCA_0: recursive comprehension axiom. WKL_0: weak König's lemma over RCA_0. Tighter localisations may be possible.

Component	Upper bound	Remarks
Tetration recursion (**??**)	$I\Sigma_1$	Primitive recursive; well-formedness in fragment of PA
Strict monotonicity of tetration (Prop. **??**)	PA	Induction on c using Σ_1-induction
Fundamental Theorem of Arithmetic	PA	Classical; Bezout / well-ordering argument
No-Tie Lemma (Lemma **??**)	PA	Elementary arithmetic and p-adic valuations
Greedy peel termination (Thm. **??**)	PA	Strong induction via strict descent of X
Hyperfactorization Theorem (Thm. **??**)	PA	Assembles the above via strong induction

Summary. Every part of the theorem is provable in Peano Arithmetic; no analytic or set-theoretic resources are required. The theorem is in this sense *elementary*: the statement and proof can be carried out entirely within the finitary fragment of arithmetic, matching its role as a foundational hinge.

B. LEAN 4 FORMALISATION STATUS AND PLAN

B.1 Current TauLib state

The relevant TauLib modules at the commit-of-record are:

Category A — formally proved theorems..

- `tetration_unbounded`, `tetration_strict_mono` in `TauLib/BookI/Coordinates/Hyperfact.lean`: strict monotonicity and unboundedness of tetration.
- `no_tie` in `TauLib/BookI/Coordinates/NoTie.lean`: Part (1) (maximum tetration height exists) and Part (2) (tower-atom injectivity).
- Supporting arithmetic results in `TauLib/BookI/Polarity/NthPrime.lean` (prime-counting, primality of `nth_prime`).

Category B — computable definitions and decision procedures..

- `hyperfact_check`, `valid_abcd_check`, `encoding_check` in `TauLib/BookI/Coordinates/Hyperfact.lean`.
- `abcd_chart`, `coord_A`, `coord_B`, `coord_C`, `coord_D`, `chart_value` in `TauLib/BookI/Coordinates/ABCD.lean`.
- `greedy_peel`, `tower_atom` in `TauLib/BookI/Coordinates/TowerAtoms.lean`.

- `spine` in `TauLib/BookI/Coordinates/NormalForm.lean` (iterated peel).

Category C — claimed but not present as theorems..

- *Hyperfactorization Theorem* (I.T04): no Lean `theorem` of this content; only `hyperfact_check` as a computable verifier.
- *ABCD injectivity*: no Lean `theorem`.
- *Shadow-equality collapse*: no Lean `theorem`.

B.2 Proof-chain sketch

We sketch how each part of Theorem **??** lifts to a Lean 4 proof over `mathlib` and TauLib.

Step 1: Tetration and tower atoms.. TauLib already defines `tetration` (in `Hyperfact.lean`) as the recursive definition (**??**), with formally proved monotonicity. `mathlib`'s `Nat.factorization` and `Nat.Primes` provide the FTA primitives. For $X \geq 2$, the largest prime divisor of X is `(Nat.factorization X).support.max' h`, where `h` is a non-emptiness proof of `(Nat.factorization X).support`; this returns a concrete `Nat` (rather than the `WithBot Nat` from the partial version `Finset.max`). The edge case $X \in \{0, 1\}$ is outside the paper's theorem scope and is handled by the precondition `hX : X` $\geq$ `2`, which also discharges `h`.

Step 2: Greedy peel.. Definition **??** is already implemented as `greedy_peel` in `TowerAtoms.lean`; the five-step procedure is purely computational. A Lean theorem `greedy_peel_admissible` asserting that the output is an admissible ABCD tuple (Proposition **??**) is a direct proof from definitions plus the existing `Nat.factorization` API.

Step 3: No-Tie Lemma.. `NoTie.lean` already proves both parts of Lemma **??** as `theorem no_tie`. This is the key rigidity ingredient and is already in place.

Step 4: Strict remainder descent.. Lemma **??** and Theorem **??** are implemented computationally in `Descent.lean` and `NormalForm.lean`; formal theorem statements `remainder_descent` (with proof by divisor arithmetic) and `spine_terminates` (proof by strong induction) are straightforward additions.

Step 5: Hyperfactorization uniqueness.. The strong-induction proof of Theorem **??** assembles Steps 2–4 with FTA (for Step 1 of the four-step proof) and cancellation in $\mathbb{N}$. All sub-ingredients are available in `mathlib` or TauLib; the main work is writing the induction proof (100 Lean lines).

Step 6: ABCD injectivity and shadow collapse.. Immediate consequences (Corollaries **??** and **??**) — approximately 20–30 Lean lines each.

Step 7: Isomorphism Theorem and FTA embedding.. Theorem **??** is a definitional equality (both sides are the same computational function `greedy_peel` applied to the same input), so the Lean theorem reduces to `rfl` once the kernel identification τ-Idx $\cong \mathbb{N}_{\geq 1}$ is imported. Corollary **??** is a specialisation lemma on the recursive ABCD chain restricted to $C = 1$.

Line-count estimate.. A realistic estimate for the orthodox Theorem (Theorem **??**) plus its corollaries is 500–700 **new Lean lines** on top of the existing TauLib modules. The breakdown: 200–300 lines for the strong-induction proof itself (four case-steps, base case at $X = 2$, each invoking pre-existing lemmas through adapter lemmas against `mathlib`); 100 lines for glue formalisations of Definition **??** and Proposition **??**; 100 lines for the corollaries (injectivity, shadow collapse, FTA embedding); and 100–200 lines for API adapters against `mathlib`'s `Nat.factorization` and the existing tetration infrastructure in TauLib. Total: $200\text{–}300 + 100 + 100 + 100\text{–}200 = 500\text{–}700$ lines. The τ-framework Theorem (Theorem **??**) reduces to the orthodox one under the kernel identification, so no additional work is needed beyond importing the α-orbit bijection theorem from Book I's Lean module. Total work: one focused formalisation sprint, estimated at three to four calendar weeks.

Commit-of-record.. The `mathlib` identifiers cited above are at the main branch as of this paper's April 2026 revision (exact commit hash to be recorded in `papers/hyperfactorization/README.md` at release); identifier names may need minor adjustment at implementation time.