HINGE 2 PREPRINT · V1.0 · APRIL 2026

The Prime Polarity Theorem

A ZFC/PA proof and the τ-framework isomorphism

Thorsten Fuchs • Anna-Sophie Fuchs

Correspondence: thorsten@panta-rhei.site

April 2026

DOI: 10.5281/zenodo.19819869

ABSTRACT

The Prime Polarity Theorem asserts that the primes admit a canonical bipolar partition $\mathbb{P} = \mathbb{P}_B \sqcup \mathbb{P}_C \sqcup \{2\}$ into classes distinguished by the quadratic character of 2 modulo p. For odd primes p,

$$p \in \mathbb{P}_B \iff \left(\tfrac{2}{p}\right) = +1 \iff p \equiv \pm 1 \pmod 8, \qquad p \in \mathbb{P}_C \iff \left(\tfrac{2}{p}\right) = -1 \iff p \equiv \pm 3 \pmod 8,$$

and 2 carries the mixed label X. Both $\mathbb{P}_B$ and $\mathbb{P}_C$ are infinite (Dirichlet, 1837) of natural density $\frac{1}{2}$ (Prime Number Theorem for arithmetic progressions; de la Vallée Poussin / Siegel–Walfisz).

We establish this theorem in two forms and prove their exact equivalence.

Orthodox form (ZFC).. A direct classical proof via Euler's criterion, the second supplementary law of quadratic reciprocity [14], and Dirichlet's theorem for residues modulo 8 (Theorem 5.3).

τ-framework form.. A structural derivation inside Category τ starting from the Chinese Remainder Theorem idempotents of the primorial ring $\mathbb{Z}/\mathrm{Prim}(n)\mathbb{Z}$, extended by the split-complex boundary ring $H_\tau = \mathbb{Z}[\mathrm{j}]/(\mathrm{j}^2 - 1)$, and read off at each prime via a localised-Legendre convention (the *spectral Legendre symbol* $\mathrm{SpecLeg}_n$; Definition 6.2) that survives the CRT-orthogonality collapse. The resulting internal bipolar classifier Label_∞ is stable at the depth each prime first enters (Theorem 6.8).

Isomorphism Theorem.. The orthodox classifier Pol and the τ-framework classifier Label_∞ agree pointwise on the primes (Theorem 7.1). This is the paper's central technical contribution: it bridges the structural derivation with the classical Gauss classification, so that downstream τ-framework constructions (boundary characters, L-functions, the spectral algebra of the lemniscate) are compatible with orthodox analytic number theory.

We also identify two candidate non-Legendre classifiers that fail: the universal–existential growth-rate criterion is vacuous (Theorem 2.1), and the bound-dependent spectral signature has asymptotically empty C-class (Theorem 3.1). A precise intrinsic-growth-rate conjecture aligned with the Legendre classification via a primorial-ladder scaled-gap inequality is stated as Conjecture 8.1, offered as an open problem in analytic number theory.

The axiomatic strength of each part of the theorem is located in Appendix A. A proof-chain sketch of a Lean 4 formalisation, targeting `mathlib`'s analytic and algebraic number-theory primitives, is given in Appendix B.

Keywords Prime polarity ; Bipolar partition ; Legendre symbol ; Kronecker symbol ; Quadratic residues ; Dirichlet's theorem ; Chinese Remainder Theorem ; Split-complex idempotents ; Category τ ; Lean 4 formalisation

MSC 2020 Mathematics Subject Classification: 11A41, 11A07, 11A15, 11N13, 18A15, 03F65, 68V20

1. INTRODUCTION

1.1 Motivation and scope

The *Panta Rhei* monograph series [10, 11, 12] develops Category τ, a categorical framework whose mathematical stratum (Books I–III) reconstructs portions of number theory, complex analysis, spectral theory, and the Langlands programme from a kernel of seven axioms, five generators, and one progression operator. A central theorem of this stratum is the *Prime Polarity Theorem*[1] (I.T05 [10]): the primes admit a canonical bipolar partition $\mathbb{P} = \mathbb{P}_B \sqcup \mathbb{P}_C \sqcup \{2\}$ with both classes infinite. This theorem is cited throughout the program, notably in the bundle-companion Master Constant paper [9], which uses the Legendre classifier established here as the orthodox image of its split-complex idempotent character; and in the framework-application papers on black-hole stability [7] and arithmetic quantum gravity [6], where the bipolar partition is the input to the boundary-character decomposition that powers the L-function construction of Book III [12].

The purpose of this paper is to establish the theorem rigorously in both an orthodox ZFC form and a τ-framework form, and to prove the pointwise equivalence of the two classifiers so obtained. The equivalence — the *Isomorphism Theorem* (Theorem 7.1) — is the bridge between the classical quadratic-residues route (Euler, Gauss, Dirichlet) and the structural τ-derivation via CRT idempotents and split-complex scalars.

1.2 Structure of the paper

We make the following three mathematical contributions:

(1) **Orthodox theorem.** The Legendre-symbol classifier Pol partitions the primes; both classes are infinite of natural density $\frac{1}{2}$ (Theorem 5.3). The proof uses Euler's criterion, the second supplementary law of quadratic reciprocity (Gauss 1801 [14]), and Dirichlet's theorem [5]. See §5; the axiomatic strength of each part is located in Appendix A.

(2) **τ-framework derivation.** The same classifier arises canonically inside Category τ from the CRT idempotents on the primorial ring, extended by the split-complex boundary ring $H_\tau = \mathbb{Z}[\mathrm{j}]/(\mathrm{j}^2 - 1)$, and read off at each prime via the spectral Legendre symbol $\mathrm{SpecLeg}_n$ (§6, Theorem 6.8).

(3) **Isomorphism.** The orthodox classifier Pol and the τ-framework classifier Label_∞ agree pointwise as functions $\mathbb{P} \to \{B, C, X\}$ (§7, Theorem 7.1).

We also establish two diagnostic results that rule out alternative classifier candidates:

(D1) **Triviality of the $\forall\exists$ criterion.** A universal-existential criterion comparing p^B with $p \uparrow\uparrow C$ on tower-atom witnesses holds for every prime, and so does not partition $\mathbb{P}$ (Theorem 2.1).

(D2) **Bound artifact of the spectral signature.** The bound-dependent spectral signature $\sigma_N(p) = (B_{\max}(p, N), C_{\max}(p, N))$, sampled from τ-index objects $X \le N$, produces a classifier whose C-class is asymptotically empty as $N \to \infty$ (Theorem 3.1).

These diagnostics (§2–§4) motivate why the Legendre-symbol classifier is the right object for the partition theorem.

Axiomatic calibration.. The classification and dichotomy parts of the theorem (Parts 1–2 of Theorem 5.3) are provable in PA; the infinitude claim (Part 3) is provable in RCA_0 via the elementary Selberg–Erdős route to Dirichlet; the density claim (Part 4) uses the Prime Number Theorem for arithmetic progressions (de la Vallée Poussin, Siegel–Walfisz) [15, 18]. Appendix A gives the locator. Lean formalisation is discussed in Appendix B.

1.3 Historical note and related work

The observation that the Legendre symbol $\left(\frac{2}{p}\right)$ partitions the primes into four residue classes modulo 8 (with $p \equiv \pm 1$ giving $+1$ and $p \equiv \pm 3$ giving -1) goes back to Euler and Gauss [14]. Dirichlet [5] established infinitude of primes in each nonzero residue class modulo a given modulus. The Chebotarev density theorem [3] generalised this to give densities for primes splitting in a fixed way in a finite Galois extension; the mod-8 case reduces to the easier Dirichlet density calculation (both yield density $\frac{1}{2}$ for each of $\mathbb{P}_B$ and $\mathbb{P}_C$). For classical accounts, see Serre [20], Apostol [1], Ireland–Rosen [17], Iwaniec–Kowalski [18].

The contribution of the τ-framework is not the classical mod-8 result but the *derivation* of the classifier from a structural-algebraic principle (the CRT idempotent + split-complex decomposition), which is what makes the classifier *earned* rather than

[1]The present paper is **Hinge 2** of the Panta Rhei four-hinge standalone-paper bundle, in recommended reading order: **Hinge 1** — the *Hyperfactorization Theorem* [8] (I.T04), which supplies the tower-atom coordinate functions used in Lemma 2.1 and Lemma 3.1; **Hinge 2** — the *Prime Polarity Theorem* treated here (I.T05), which uses those coordinates to classify primes via the Legendre symbol $(2/p)$; **Hinge 3** — the *Master Constant* ι_τ paper [9], which lifts the character χ derived here to a split-complex idempotent character $\widetilde{\chi}$ and identifies $\iota_\tau = 2/(\pi + e)$ as the lemniscate crossing-germ scalar; and **Hinge 4** — the *Split-Complex Boundary Algebra* paper [13], 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 hinges (notably III.T19 Critical Line) are established in later books of the series.

postulated inside Category τ. The Isomorphism Theorem below says that the structural derivation recovers exactly the classical Legendre classification.

2. DIAGNOSTIC: TRIVIALITY OF THE ∀∃ DEFINITION

The most natural candidate definition of B-dominance and C-dominance compares pure-power tower atoms p^B with tetration tower atoms $p \uparrow\uparrow C$ inside τ-Idx.

2.1 The candidate ∀∃ criterion

For a prime $p \in \mathbb{P}$, consider the *universal-existential criterion*:

- p is **B-dominant** if for every $C \geq 2$, there exist $B \geq 2$ and $X \in \tau$-Idx with $\mathrm{coord}_A(X) = p$, $\mathrm{coord}_B(X) = B$, $\mathrm{coord}_C(X) = 1$, and
$$p^B > p \uparrow\uparrow C; \tag{1}$$
- p is **C-dominant** if for every $B \geq 2$, there exist $C \geq 2$ and $X \in \tau$-Idx with $\mathrm{coord}_A(X) = p$, $\mathrm{coord}_C(X) = C$, $\mathrm{coord}_B(X) = 1$, and
$$p \uparrow\uparrow C > p^B. \tag{2}$$

An initially plausible statement of the Prime Polarity Theorem would read: $\mathbb{P} = \mathbb{P}_B \sqcup \mathbb{P}_C$ under this criterion with both classes infinite. The next subsection shows that such a statement is vacuous.

2.2 Both conditions are trivially true for every prime

The growth-rate separation theorem, established in Book I Chapter 12 and formalised in Lean as

$$\texttt{growth\_rate\_separation} : \forall a \geq 2,\ \forall B,\ \exists C,\ a \uparrow\uparrow C > a^B \qquad (\text{TauLib: }\texttt{Spectral.lean:77}) \tag{3}$$

and its dual

$$\texttt{b\_beats\_c} : \forall a \geq 2,\ \forall C,\ \exists B,\ a^B > a \uparrow\uparrow C \qquad (\text{TauLib: }\texttt{Polarity.lean:239}) \tag{4}$$

directly imply the following.

Lemma 2.1 (Vacuity of the ∀∃ criterion). *For every prime $p \geq 2$:*

(i) the B-dominance condition (1) *holds, and*

(ii) the C-dominance condition (2) *holds.*

In particular, the definition does not partition the primes.

Proof. We give an argument that uses only the coordinate functions' definition via greedy peel-off and the Lean-proved growth theorems. The argument is a *specialisation* of Hyperfactorization (I.T04) to pure-power inputs: it reconstructs locally the $D = 1$ case of the full decomposition without invoking the theorem globally. The facts used about $(\mathrm{coord}_A, \mathrm{coord}_B, \mathrm{coord}_C)$ in the pure-power regime reduce to the elementary observation that $(p \uparrow\uparrow c)^b$ divides p^B only if $t_c \mid B$ (where $t_c = \log_p(p \uparrow\uparrow c)$), which is standard divisor arithmetic and does not require I.T04 as a prerequisite.

(i) B-dominance of every prime. Fix $C \geq 2$. By `b_beats_c`, there exists B_0 with $p^{B_0} > p \uparrow\uparrow C$. Choose $B \geq \max(B_0 + p, 2)$ with $\gcd(B, p) = 1$; such a B exists because $\{B_0 + p, B_0 + p + 1, \ldots\}$ contains infinitely many integers coprime to p. Then $p^B = p^{B-B_0} \cdot p^{B_0} > p \uparrow\uparrow C$. Consider the pure power $X = p^B$:

- $\mathrm{coord}_A(X) = p$: by definition (largest prime divisor of p^B is p).
- $\mathrm{coord}_C(X) = 1$: by definition of the greedy peel, $\mathrm{coord}_C(X)$ is the largest c with $(p\uparrow\uparrow c)^b \mid X$ and quotient $X/(p\uparrow\uparrow c)^b$ having all prime factors $< p$. For $c \geq 2$, $p \uparrow\uparrow c = p^{t_c}$ with $t_c := \log_p(p \uparrow\uparrow c) \geq p$; hence $(p \uparrow\uparrow c)^b = p^{b \cdot t_c}$ divides p^B only if $t_c \mid B$, which requires $p \mid B$ and contradicts $\gcd(B, p) = 1$. Therefore $c = 1$ is forced.
- $\mathrm{coord}_B(X) = B$: with $c = 1$ fixed, $(p \uparrow\uparrow 1)^b \cdot D = p^b \cdot D = p^B$ with D having all prime factors $< p$ requires $D = 1$ (for $p = 2$) or D a power of smaller primes (for $p > 2$, but here X is a pure power of p so $D = 1$). Thus $b = B$.

Hence $X = p^B$ witnesses (1).

(ii) C-dominance of every prime. Fix $B \geq 2$. By `growth_rate_separation`, there exists C with $p \uparrow\uparrow C > p^B$. Any such witness has $C \geq 2$ since $p \uparrow\uparrow 1 = p \leq p^B$. Consider the tower $X = p \uparrow\uparrow C$:

- $\mathrm{coord}_A(X) = p$: the only prime factor of $p \uparrow\uparrow C$ is p.
- $\mathrm{coord}_C(X) = C$: by definition, greedy peel picks the largest c with $(p \uparrow\uparrow c)^b \cdot D = p \uparrow\uparrow C$ and D's prime factors $< p$. For $c = C, b = 1, D = 1$: valid. For $c > C$, $(p \uparrow\uparrow c)^1 > p \uparrow\uparrow C$: invalid. Hence $c = C$ is maximal.
- $\mathrm{coord}_B(X) = 1$: with $c = C$, $(p \uparrow\uparrow C)^b \cdot D = p \uparrow\uparrow C$ requires $b = 1$ (any $b \geq 2$ would give $(p \uparrow\uparrow C)^b > p \uparrow\uparrow C$).

Hence $X = p \uparrow\uparrow C$ witnesses (2). □

Consequence.. Under the ∀∃ criterion, $\mathbb{P}_B = \mathbb{P}_C = \mathbb{P}$: the two sets coincide with $\mathbb{P}$ rather than partitioning it. A partition theorem therefore requires a different classifier. The next section examines the natural bound-dependent alternative and shows that it too falls short.

3. THE BOUND-DEPENDENT SPECTRAL SIGNATURE IS NOT ENOUGH

3.1 Spectral signature $\sigma_N(p)$

The *spectral signature at bound N* is defined [10] as

$$\sigma_N(p) \;:=\; (B_{\max}(p, N),\; C_{\max}(p, N)), \tag{5}$$

where, writing $S_p(N) := \{X \in \mathbb{N} :\; 2 \leq X \leq N,\; \mathrm{coord}_A(X) = p\}$,

$$B_{\max}(p, N) \;:=\; \max_{X \in S_p(N)} \mathrm{coord}_B(X), \tag{6}$$

$$C_{\max}(p, N) \;:=\; \max_{X \in S_p(N)} \mathrm{coord}_C(X). \tag{7}$$

(Book I Def. I.D19e; Lean: `spectral_sig` in `TauLib/BookI/Polarity/Spectral.lean`, lines 35–58.) The coordinate maps coord_A, coord_B, coord_C come from the *Hyperfactorization Theorem* (I.T04, [10] Chapter 24; bundle-companion paper [8], Hinge 1): every integer $X \geq 2$ has a unique decomposition $X = T(A, B, C) \cdot D$ with A prime, $B, C \geq 1$, $D \geq 1$, all prime factors of D strictly less than A, and $T(A, B, C) := (A \uparrow\uparrow C)^B$ the "tower atom."

The bound-dependent polarity at bound N is

$$\mathrm{pol\_at}(p, N) \;:=\; [\![B_{\max}(p, N) > C_{\max}(p, N)]\!], \tag{8}$$

i.e. p is *B-dominant at bound N* iff $B_{\max} > C_{\max}$ at that bound. One might hope to pass to an intrinsic classifier via asymptotic stabilisation as $N \to \infty$.

3.2 The asymptotic classifier collapses to the B-class

Write $b_N(p) := B_{\max}(p, N)$ and $c_N(p) := C_{\max}(p, N)$. We show that for every prime p, $b_N(p) > c_N(p)$ for all sufficiently large N, so the asymptotic classifier assigns every prime to the B-class. Hence the asymptotic C-class is empty, and the partition-and-infinitude theorem cannot rest on this classifier.

Lemma 3.1 (Asymptotic vacuity of spectral-signature C-class)**.** *For every prime $p \geq 2$ there exists $N_0(p) < \infty$ such that $b_N(p) > c_N(p)$ for all $N \geq N_0(p)$. Consequently, $\{p \in \mathbb{P} :\; b_N(p) \leq c_N(p)$ for all sufficiently large $N\} = \emptyset$.*

Proof. We first show that for every prime $p \geq 2$, $b_N(p)$ grows like $\log_p N$ while $c_N(p) \leq \log_p^* N + O(1)$, so that $b_N(p) > c_N(p)$ for N beyond a finite threshold.

Lower bound on $b_N(p)$ via coprime-exponent witnesses. For $p = 2$, consider pure powers $X = 2^k$ with k *odd*. By the computation in the proof of Lemma 2.1(i) (greedy-peel analysis under $\gcd(k, 2) = 1$), such X have $\mathrm{coord}_A(X) = 2$, $\mathrm{coord}_C(X) = 1$, and $\mathrm{coord}_B(X) = k$. Every odd $k \leq \lfloor \log_2 N \rfloor$ yields an $X \leq N$ with $\mathrm{coord}_B = k$; hence

$$b_N(2) \;\geq\; \max\{k \text{ odd} :\; k \leq \lfloor \log_2 N \rfloor\} \;\geq\; \lfloor \log_2 N \rfloor - 1. \tag{9}$$

For odd prime $p \geq 3$, the same argument applies with any k coprime to p: such k exist in every length-p window, so $k = \lfloor \log_p N \rfloor$ or $k = \lfloor \log_p N \rfloor - 1$ is coprime to p, giving $b_N(p) \geq \lfloor \log_p N \rfloor - 1$.

Upper bound on $c_N(p)$ via tower containment. For any $X \in S_p(N)$, $\mathrm{coord}_C(X) \leq \max\{c : p \uparrow\uparrow c \leq N\}$ (since $\mathrm{coord}_C(X) = c$ requires $p \uparrow\uparrow c \mid X \leq N$, hence $p \uparrow\uparrow c \leq N$). The quantity $\max\{c : p \uparrow\uparrow c \leq N\}$ grows like the p-iterated logarithm $\log_p^* N$; concretely, for $p = 2$ it equals 4 for all $N \in [2^4, 2^{16} - 1]$ and 5 for $N \in [2^{16}, 2^{65536} - 1]$, etc. Hence

$$c_N(p) \;\leq\; \log_p^* N + O(1). \tag{10}$$

Crossing point. Combining (9)–(10): for every prime p, $b_N(p) - c_N(p) \geq \lfloor \log_p N \rfloor - \log_p^* N - O(1) \to +\infty$ as $N \to \infty$. A concrete threshold for $p = 2$ is $N_0(2) = 2^5 = 32$: at $N = 32$, $b_{32}(2) \geq 4$ (from $X = 2^5 = 32$, odd $k = 5$), $c_{32}(2) \leq 3$ (since $2 \uparrow\uparrow 4 = 65536 > 32$). For odd prime p, $N_0(p) = p^p$ suffices: at $N = p^p$, $b_N(p) \geq p - 1$ (from $X = p^k$, odd $k = p$ or $p - 1$ coprime to p) while $c_N(p) \leq 2$ (since $p \uparrow\uparrow 3 = p^{p^p} \gg p^p$).

The second statement follows: for each prime p, p belongs to the B-class at bound N for all $N \geq N_0(p)$; hence the asymptotic C-class (primes with $c_N \geq b_N$ eventually) is empty. □

Interpretation.. The growth-rate separation says tetration eventually dominates exponentiation *vertically* (at fixed tower height, arbitrarily many exponentiation steps are bounded by one more tetration step). But the spectral signature samples *horizontally* (all $X \leq N$): the number of accessible coord_B-levels at bound N grows like $\log N$, whereas the number of accessible coord_C-levels grows like $\log^* N$. So the horizontal sampling is systematically biased toward B, and the asymptotic classifier places every prime in the B-class.

This is consistent with the Lean empirical data: at any finite N there is a crossover prime below which the prime is B-dominant and above which it is C-dominant, but the crossover moves to $+\infty$ as N grows.

4. EMPIRICAL DATA: LEAN AT $N = 1000$

To make the bound-artifact concrete, we ran the Lean `polarity_chi` function (TauLib commit-of-record, `BookI/Polarity/Polarity.lean`, lines 134–138) at bound $N = 1000$ for the first fifteen primes. The computed polarity (with the Lean convention $\chi = -1 \Rightarrow B$-dominant, $\chi = +1 \Rightarrow C$-dominant at bound N) is reported in Table 1.

Scaling of the crossover prime.. The analytical prediction $p^*(N) \sim \lceil \sqrt{N} \rceil$ extends uniformly: at $N = 10^6$ the crossover moves to $p^* \geq 1009$; at $N = 10^9$, $p^* \geq 31627$; and $p^*(N) \to +\infty$ as $N \to \infty$ (Lemma 3.1). The bound artifact is therefore a *genuinely divergent* deviation from Label_∞, not a finite-size quirk.

What the data show.. The bound-dependent polarity flips at $p^*(N) = \lceil \sqrt{N} \rceil$: primes $p < p^*(N)$ satisfy $p^2 \leq N$ so $B_{\max} \geq 2 > C_{\max} = 1$, making them B-dominant at bound N; primes $p \geq p^*(N)$ satisfy $p^2 > N$ so $B_{\max} = C_{\max} = 1$, making them C-dominant at bound N (by the strict-inequality convention). At $N = 1000$ the crossover is at $p^* = 37$; at $N = 10000$ the crossover has moved to $p^* \geq 101$ (the first prime with $p^2 > 10000$). Concretely: the prime $p = 41$ classifies as C-dominant at $N = 1000$ but B-dominant at $N = 10000$ — its bound-dependent polarity is not a property of 41 but of the pair $(41, N)$.

The Legendre classifier $\mathrm{Label}_\infty(p)$ does not track this bound artifact: it is N-independent by construction. The two classifiers disagree on a large fraction of primes (7 of 14 odd primes in the $N = 1000$ sample; 8 of 14 in the $N = 10000$ sample) and have entirely different asymptotic behaviour (Legendre: density $\frac{1}{2}$ each class; bound-dependent: density 0 for the C-class as $N \to \infty$, by Lemma 3.1).

The paper accordingly adopts Label_∞ as the canonical polarity. The next section states and proves the theorem in purely orthodox form; §6 derives the same classifier structurally inside Category τ.

5. THE ORTHODOX THEOREM: ZFC PROOF

We work in ZFC. The partition and dichotomy parts are elementary and provable already in PA; the infinitude and density parts rely on Dirichlet's theorem on primes in arithmetic progressions (and its density refinement), which is classical analytic number theory. A reverse-mathematical locator of which parts need which fragments appears in §A.

Table 1. Bound-dependent polarity $\chi(p, N)$ versus the Legendre-classifier $\mathrm{Label}_\infty(p)$. The $\chi_{\mathrm{Lean}}(p, 1000)$ column is a verbatim Lean `polarity_chi` run on `TauLib/BookI/Polarity/Polarity.lean` (commit-of-record, April 2026). The $\chi_{\mathrm{pred}}(p, 10000)$ column is the analytical prediction from Theorem 3.1: a prime is B-dominant at bound N iff $p^2 \le N$ (concretely, iff $p < p^*(N) := \lceil \sqrt{N} \rceil$). The analytical prediction is not independently Lean-verified at $N = 10000$ (direct `spectral_sig` evaluation at that bound takes more than 7 wall-clock minutes at commit-of-record; a future Lean plan milestone is to Lean-verify these rows by short-circuiting the scan when $p^2 > N$). A prime has $\mathrm{Label}_\infty(p) = B$ iff $p \equiv \pm 1 \pmod 8$; otherwise $\mathrm{Label}_\infty(p) = C$ (for $p = 2$: X). Disagreements with Label_∞ highlighted with †. The crossover prime p^* moves from 37 at $N = 1000$ to ≥ 101 at $N = 10000$, concretely exhibiting the bound-artifact behaviour of Theorem 3.1.

p	p mod 8	$\mathrm{Label}_\infty(p)$	$\chi_{\mathrm{Lean}}(p, 1000)$	$\chi_{\mathrm{pred}}(p, 10000)$	Matches Label_∞?
2	2	X	-1 (B)	-1 (B)	–
3	3	C	-1 (B)	-1 (B)	††
5	5	C	-1 (B)	-1 (B)	††
7	7	B	-1 (B)	-1 (B)	□□
11	3	C	-1 (B)	-1 (B)	††
13	5	C	-1 (B)	-1 (B)	††
17	1	B	-1 (B)	-1 (B)	□□
19	3	C	-1 (B)	-1 (B)	††
23	7	B	-1 (B)	-1 (B)	□□
29	5	C	-1 (B)	-1 (B)	††
31	7	B	-1 (B)	-1 (B)	□□
37	5	C	$+1$ (C)	-1 (B)	□†
41	1	B	$+1$ (C)	-1 (B)	†□
43	3	C	$+1$ (C)	-1 (B)	□†
47	7	B	$+1$ (C)	-1 (B)	†□

5.1 Definition and main theorem

Definition 5.1 (Orthodox prime polarity). *Define* $\mathrm{Pol} : \mathbb{P} \to \{B, C, X\}$ *by*

$$\mathrm{Pol}(p) \;:=\; \begin{cases} X, & p = 2, \\ B, & p > 2 \text{ and } \left(\frac{2}{p}\right) = +1, \\ C, & p > 2 \text{ and } \left(\frac{2}{p}\right) = -1, \end{cases} \tag{11}$$

where $\left(\frac{2}{p}\right)$ *is the Legendre symbol. Set* $\mathbb{P}_B := \mathrm{Pol}^{-1}(B)$, $\mathbb{P}_C := \mathrm{Pol}^{-1}(C)$, $\mathbb{P}_X := \mathrm{Pol}^{-1}(X) = \{2\}$.

Remark 5.2 (Mod-8 explicit form). By the *second supplementary law of quadratic reciprocity* (Gauss, *Disquisitiones Arithmeticae* §112, 1801 [14]; modern account in Ireland–Rosen [17, §5.2]),

$$\left(\frac{2}{p}\right) \;=\; (-1)^{(p^2-1)/8} \;=\; \begin{cases} +1 & \text{if } p \equiv \pm 1 \pmod 8, \\ -1 & \text{if } p \equiv \pm 3 \pmod 8, \end{cases} \tag{12}$$

for every odd prime p. Thus $\mathrm{Pol}(p) = B$ iff $p \equiv \pm 1 \pmod 8$, and $\mathrm{Pol}(p) = C$ iff $p \equiv \pm 3 \pmod 8$. We note that the *first* supplementary law, governing $\left(\frac{-1}{p}\right)$, is not used in this paper; only the second supplement is.

Theorem 5.3 (Prime Polarity Theorem, orthodox form). *(1)* (Partition.) $\mathbb{P} = \mathbb{P}_B \sqcup \mathbb{P}_C \sqcup \mathbb{P}_X$, *and the three classes are pairwise disjoint.*

(2) (Dichotomy for odd primes.) *For every odd prime* p, $\mathrm{Pol}(p) \in \{B, C\}$ *and exactly one of these holds.*

(3) (Infinitude.) *Both* $\mathbb{P}_B$ *and* $\mathbb{P}_C$ *are infinite.*

(4) (Density.) *Each of* $\mathbb{P}_B, \mathbb{P}_C$ *has natural density* $\frac{1}{2}$ *in the primes.*

Proof. Parts (1) and (2) are immediate from Definition 5.1 and Remark 5.2: for odd p, $\left(\frac{2}{p}\right) \in \{+1, -1\}$ and the cases $p \equiv \pm 1$ vs. $p \equiv \pm 3 \pmod 8$ are mutually exclusive and exhaustive among odd residues modulo 8.

Part (3): by Dirichlet's theorem on primes in arithmetic progressions [5, 20], each of the four arithmetic progressions $1, 3, 5, 7 \pmod 8$ contains infinitely many primes. In particular, $\{p \equiv 1 \pmod 8\} \cup \{p \equiv 7 \pmod 8\} \subseteq \mathbb{P}_B$ is infinite, and $\{p \equiv 3 \pmod 8\} \cup \{p \equiv 5 \pmod 8\} \subseteq \mathbb{P}_C$ is infinite.

Part (4): we claim natural density $\frac{1}{2}$ for each of $\mathbb{P}_B, \mathbb{P}_C$. Two routes are available.

Route A (Dirichlet density, 1837 methods). Dirichlet's 1837 method (non-vanishing of $L(1, \chi)$ for non-principal characters [5]), combined with the standard Dirichlet-density calculation via orthogonality of characters [20, §VI.4], yields *Dirichlet density* $\frac{1}{\varphi(8)} = \frac{1}{4}$ for each of the four reduced residue classes $1, 3, 5, 7 \pmod 8$. Summing over the two classes comprising $\mathbb{P}_B$ (resp. $\mathbb{P}_C$) gives Dirichlet density $\frac{1}{2}$ for each. The infinitude statement of Theorem 5.3(3) is the qualitative corollary.

Route B (natural density, de la Vallée Poussin / Siegel–Walfisz). The *natural density* $\frac{1}{2}$ follows from the Prime Number Theorem for arithmetic progressions (PNT for AP), due to de la Vallée Poussin (1896) for the qualitative statement and Siegel–Walfisz (1936) for the effective uniformity in q: for coprime a, q, $\pi(x; q, a) \sim \frac{1}{\varphi(q)} \pi(x)$ as $x \to \infty$ [15, §22.14],[18, Ch. 5]. Applying with $q = 8$ and summing over the two residue classes of each polarity yields $|\mathbb{P}_B \cap [2, x]| \sim |\mathbb{P}_C \cap [2, x]| \sim \frac{1}{2}\pi(x)$, i.e. natural density $\frac{1}{2}$ each.

The two routes coincide whenever both densities exist (as they do here); Route A is logically weaker (uses only $L(1, \chi) \neq 0$), Route B is sharper (uses the full PNT for AP, with Siegel–Walfisz effective error bounds available if wanted). Either suffices for Part (4). ☐

Remark 5.4 (Alternative density proof)**.** One can also derive part (4) from the Chebotarev density theorem [3] applied to the Galois extension $\mathbb{Q}(\zeta_8)/\mathbb{Q}$ (whose Galois group is $(\mathbb{Z}/8\mathbb{Z})^\times \cong (\mathbb{Z}/2\mathbb{Z})^2$), giving density $\frac{1}{2}$ to the two cosets of the kernel of the character $\left(\frac{2}{\cdot}\right)$. The two approaches are equivalent for this abelian case.

5.2 The classifier is computable in $O(\log p)$

Proposition 5.5 (Computability of Pol)**.** *Given an odd prime p, $\mathrm{Pol}(p)$ can be computed in $O(\log p)$ multiplications modulo p via Euler's criterion: $\mathrm{Pol}(p) = B$ iff $2^{(p-1)/2} \equiv +1 \pmod p$, and $\mathrm{Pol}(p) = C$ iff $2^{(p-1)/2} \equiv -1 \pmod p$. In bit-complexity terms, this is $\tilde{O}((\log p)^2)$ under any standard integer-multiplication algorithm (schoolbook, Karatsuba, or the Harvey–van der Hoeven $O(n \log n)$ bit-complexity [16]; the polylogarithmic factor is absorbed in $\tilde{O}$). The simpler mod-8 test (Remark 5.2) avoids modular exponentiation entirely and decides $\mathrm{Pol}(p)$ in $O(\log p)$ bit-operations.*

Proof. Euler's criterion [17, §5.2] states that for an odd prime p and a coprime to p, $a^{(p-1)/2} \equiv \left(\frac{a}{p}\right) \pmod p$. Applying with $a = 2$ and reducing gives the multiplicative test. The cost is dominated by modular exponentiation: $O(\log p)$ multiplications modulo p via repeated squaring, each of bit-complexity $O(M(\log p))$ where $M(n)$ is the cost of n-bit integer multiplication, giving total $O(\log p \cdot M(\log p)) = \tilde{O}((\log p)^2)$. The mod-8 test reads the three lowest bits of p and decides in $O(\log p)$ bit-operations (the dominant cost is reading the input). ☐

Definition 5.1 and Theorem 5.3 are entirely self-contained in ZFC (with the axiomatic-strength refinement analysed in Appendix A). They mention no Category τ primitive, no τ-index, no tower atom. We now turn to the τ-framework derivation.

6. THE τ-FRAMEWORK DERIVATION

This section gives a self-contained derivation of the Legendre classifier inside Category τ from CRT idempotents on the primorial ring extended by split-complex scalars. Our source is Book III Chapter 18 [12]. Figure 1 summarises the derivation pipeline alongside the orthodox route; the two chains meet at the Isomorphism Theorem (Theorem 7.1).

6.1 The primorial ring and its CRT idempotents

Let $p_1 < p_2 < p_3 < \cdots$ enumerate the primes. The *primorial at depth n* is

$$\mathrm{Prim}(n) := \prod_{i=1}^{n} p_i, \tag{13}$$

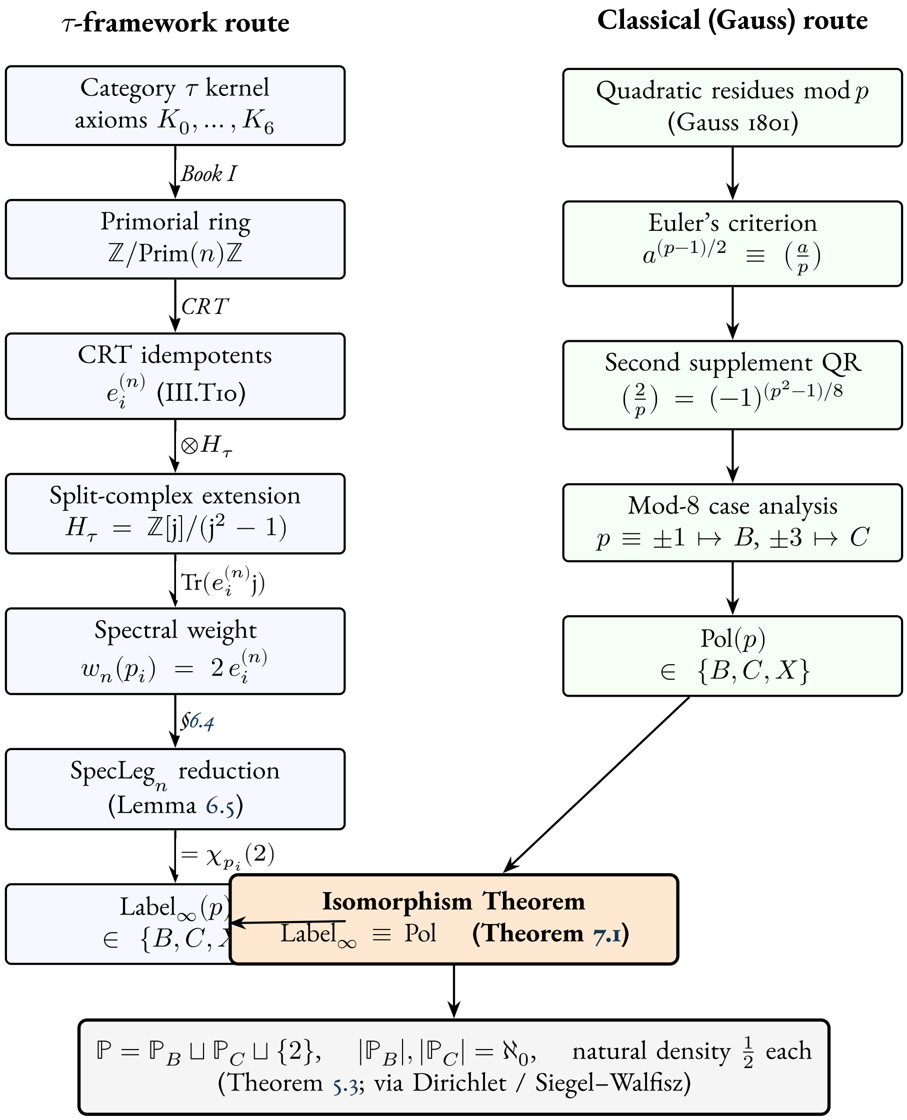

Figure 1. Pipeline of the two derivations that this paper establishes as equivalent. The left column traces the τ-framework derivation of the bipolar classifier Label_∞ through CRT idempotents and the split-complex boundary ring H_τ; the right column traces the classical derivation of Pol via the second supplement of quadratic reciprocity. The Isomorphism Theorem (Theorem 7.1) identifies the two classifiers pointwise, yielding the partition with both classes infinite of natural density $\frac{1}{2}$.

and the *primorial ring at depth* n is $\mathbb{Z}/\mathrm{Prim}(n)\mathbb{Z}$. By the Chinese Remainder Theorem, the natural map

$$\Phi_n \ : \ \mathbb{Z}/\mathrm{Prim}(n)\mathbb{Z} \ \xrightarrow{\sim} \ \prod_{i=1}^{n} \mathbb{Z}/p_i\mathbb{Z} \tag{14}$$

is a ring isomorphism (classical; reproduced at the categorical level as III.T10 [12]). Let $e_i^{(n)} \in \mathbb{Z}/\mathrm{Prim}(n)\mathbb{Z}$ be the image under Φ_n^{-1} of the idempotent $(\delta_{i1}, \dots, \delta_{in})$. Explicitly,

$$e_i^{(n)} \ \equiv \ \begin{cases} 1 & \pmod{p_i} \\ 0 & \pmod{p_k} \text{ for } k \neq i, k \leq n \end{cases} \tag{15}$$

and $\sum_{i=1}^{n} e_i^{(n)} \equiv 1 \pmod{\mathrm{Prim}(n)}$.

6.2 The split-complex boundary ring

The *split-complex boundary ring* is

$$H_\tau \ := \ \mathbb{Z}[\mathrm{j}]/(\mathrm{j}^2 - 1), \tag{16}$$

the ring of pairs $a + b\mathrm{j}$ with $a, b \in \mathbb{Z}$ and $\mathrm{j}^2 = +1$. Its two canonical orthogonal idempotents are

$$e_+ \ := \ \tfrac{1}{2}(1 + \mathrm{j}), \qquad e_- \ := \ \tfrac{1}{2}(1 - \mathrm{j}), \tag{17}$$

satisfying $e_+ + e_- = 1$, $e_+ \cdot e_- = 0$, $e_+^2 = e_+$, $e_-^2 = e_-$. Inside $H_\tau \otimes \mathbb{Q}$ this gives a direct-sum decomposition $H_\tau \otimes \mathbb{Q} \cong \mathbb{Q}e_+ \oplus \mathbb{Q}e_-$.

Remark 6.1 (Why split-complex, not complex)**.** The choice of $\mathrm{j}^2 = +1$ rather than $i^2 = -1$ is forced by the kernel axioms of Category τ: Book I Chapter 40 [10, Part X] shows that the canonical scalar extension earned by the kernel is split-complex, corresponding to hyperbolic (not elliptic) signature on the lemniscate boundary.

At the level of this paper, the scalar choice does not affect the prime-polarity classification. Both $\mathbb{Z}[\mathrm{j}]/(\mathrm{j}^2-1)$ and $\mathbb{Z}[i]/(i^2+1)$ admit orthogonal-idempotent decompositions upon tensoring with $\mathbb{Q}$ or $\mathbb{Z}[\frac{1}{2}]$; both give rise, via the same CRT-orthogonality argument, to a classifier that reads off $\left(\frac{2}{p}\right)$ at each prime. The Legendre symbol classification of primes is therefore *invariant under the scalar choice*; the split-complex choice only matters at the downstream step where the classifier is extended to a *boundary character* on the full ring H_τ, where the idempotent decomposition $H_\tau \otimes \mathbb{Q} \cong \mathbb{Q}e_+ \oplus \mathbb{Q}e_-$ produces the two-lobe lemniscate structure [12, III.P08]. Theorem 7.1 below concerns the prime classification; the scalar choice first becomes material in the analytic continuation of the classifier to the full boundary algebra.

6.3 The spectral weight and the internal bipolar classifier

The *spectral weight* at depth n of the prime p_i is [12, III.D23]

$$w_n(p_i) \ := \ 2\, e_i^{(n)} \ \in \ \mathbb{Z}/\mathrm{Prim}(n)\mathbb{Z}. \tag{18}$$

The factor of 2 is not conventional; it arises from a specific construction inside the split-complex boundary ring, which we now sketch. On $H_\tau = \mathbb{Z}[\mathrm{j}]/(\mathrm{j}^2 - 1)$ the *reduced trace* is the linear map $\mathrm{Tr} : H_\tau \to \mathbb{Z}$ given by $\mathrm{Tr}(a + b\mathrm{j}) := 2b$ (the Galois trace with respect to the involution $\mathrm{j} \mapsto -\mathrm{j}$; Book III [12, Def. III.D23] gives the full derivation).

Extending the CRT idempotents to $H_\tau \otimes \mathbb{Z}/\mathrm{Prim}(n)\mathbb{Z}$, each CRT idempotent $e_i^{(n)}$ admits a further decomposition via the split-complex idempotents e_+, e_-:

$$e_i^{(n)} \ = \ e_+ e_i^{(n)} + e_- e_i^{(n)} \ = \ \tfrac{1}{2}(e_i^{(n)} + e_i^{(n)}\mathrm{j}) + \tfrac{1}{2}(e_i^{(n)} - e_i^{(n)}\mathrm{j}). \tag{19}$$

Applying the reduced trace to the j-component of $e_i^{(n)}\mathrm{j}$ yields $\mathrm{Tr}(e_i^{(n)}\mathrm{j}) = 2e_i^{(n)}$, which is the spectral weight $w_n(p_i)$. In words: $w_n(p_i)$ measures the *asymmetry* between the e_+- and e_--projections of the CRT idempotent $e_i^{(n)}$, pulled through the reduced trace back into $\mathbb{Z}/\mathrm{Prim}(n)\mathbb{Z}$. The split-complex scalar choice (Remark 6.1) enters *here*: without $\mathrm{j}^2 = +1$ there is no e_+/e_- decomposition, and no lobe-asymmetry to measure.

The *internal bipolar classifier* at depth n is the function [12, III.D23]

$$\mathrm{Label}_n : \{p_1, \dots, p_n\} \to \{B, C, X\} \tag{20}$$

defined, for $p_i > 2$, by reading off the sign of the Jacobi symbol of the spectral weight modulo the primorial:

$$\mathrm{Label}_n(p_i) \;:=\; \begin{cases} B & \text{if } \left(\dfrac{w_n(p_i)}{\mathrm{Prim}(n)}\right) = +1, \\ C & \text{if } \left(\dfrac{w_n(p_i)}{\mathrm{Prim}(n)}\right) = -1, \end{cases} \tag{21}$$

and $\mathrm{Label}_n(2) := X$ by convention (since $\mathbb{Z}/2\mathbb{Z}$ admits no nontrivial quadratic character). Here $\left(\frac{\cdot}{\mathrm{Prim}(n)}\right)$ is the Jacobi symbol.

6.4 The CRT-local reduction to a single Legendre symbol

The key structural identity is that the spectral weight $w_n(p_i)$ carries a single piece of local number-theoretic information — the quadratic character of 2 modulo p_i — and that this is its only nonvanishing local Legendre factor under the CRT decomposition. Stating the identity rigorously requires care with the Jacobi-symbol convention, for reasons we spell out below.

Why the classical Jacobi symbol is not what we want.. Under the classical extension of the Jacobi symbol (Ireland–Rosen [17, §5.2]), $\left(\frac{a}{p}\right) := 0$ whenever $p \mid a$, and the full symbol factors as

$$\left(\frac{a}{p_1^{e_1} \cdots p_n^{e_n}}\right) \;=\; \prod_{k=1}^{n} \left(\frac{a}{p_k}\right)^{e_k}. \tag{22}$$

For the spectral weight $w_n(p_i) = 2e_i^{(n)}$ one has $w_n(p_i) \equiv 0 \pmod{p_k}$ for all $k \neq i$, $k \leq n$ (CRT orthogonality), which under the classical convention makes

$$\left(\frac{w_n(p_i)}{\mathrm{Prim}(n)}\right)_{\text{classical}} \;=\; \prod_{k=1}^{n} \left(\frac{w_n(p_i)}{p_k}\right)_{\text{classical}} \;=\; 0 \qquad (n \geq 2) \tag{23}$$

since every factor with $k \neq i$ is 0 and a single zero factor annihilates the product. The classical Jacobi symbol therefore does not transmit the local residuosity information at p_i. A localised convention — keeping only the indices where the residue is nonzero — is needed.

The spectral Legendre convention.. For each prime $p \geq 2$ and integer a, write $\chi_p(a)$ for the local quadratic character:

$$\chi_p(a) \;:=\; \begin{cases} \left(\frac{a}{p}\right) & p \geq 3 \text{ (the Legendre symbol modulo } p\text{)}, \\ +1 & p = 2 \text{ and } a \equiv \pm 1 \pmod 8, \\ -1 & p = 2 \text{ and } a \equiv \pm 3 \pmod 8, \\ 0 & p = 2 \text{ and } a \text{ even}, \end{cases} \tag{24}$$

where the $p = 2$ case follows the *Kronecker symbol* convention of Cohen [4, §1.4.2]; at $p \geq 3$ the classical Legendre symbol satisfies $\left(\frac{a}{p}\right) = 0$ when $p \mid a$. This is a total function $\chi_p : \mathbb{Z} \to \{+1, -1, 0\}$ at every prime p.

Definition 6.2 (Spectral Legendre symbol)**.** *Given $a \in \mathbb{Z}/\mathrm{Prim}(n)\mathbb{Z}$ with CRT-components $(a_1, \dots, a_n) := \Phi_n(a) \in \prod_{k=1}^{n} \mathbb{Z}/p_k\mathbb{Z}$, the* spectral Legendre symbol *of a at depth n is*

$$\mathrm{SpecLeg}_n(a) \;:=\; \prod_{k:\, p_k \nmid a} \chi_{p_k}(a_k), \tag{25}$$

the product taken only over indices k for which $p_k \nmid a$ (equivalently, $a_k \neq 0 \bmod p_k$). If the index set is empty, $\mathrm{SpecLeg}_n(a) := 1$ (empty-product convention). This is well-defined for all a, because each factor $\chi_{p_k}(a_k)$ with $p_k \nmid a$ is in $\{+1, -1\}$ by (24).

Remark 6.3 (SpecLeg as a Kronecker refinement). The Kronecker symbol $\left(\frac{a}{m}\right)$ (Cohen [4, §1.4.2]) extends the Jacobi symbol to all $m \in \mathbb{Z}_{>0}$ (not just odd) by using $\chi_2(a)$ as defined in (24) for the $p = 2$ factor. Under Kronecker, one still has $\left(\frac{a}{m}\right) = 0$ whenever $\gcd(a, m) > 1$, because the local factor at a prime p dividing both a and m is 0 and annihilates the product. The spectral Legendre symbol $\mathrm{SpecLeg}_n$ differs from Kronecker *only* in that it *skips* those zero factors rather than absorbing them: when $\gcd(a, \mathrm{Prim}(n)) = 1$, $\mathrm{SpecLeg}_n(a) = \left(\frac{a}{\mathrm{Prim}(n)}\right)_{\mathrm{Kronecker}}$; when $\gcd(a, \mathrm{Prim}(n)) > 1$, Kronecker returns 0 while $\mathrm{SpecLeg}_n$ returns the product of local factors at primes where a is nonzero. This localisation is the natural reading of a residuosity structure from a CRT-decomposed element.

Example. Take $n = 3$ ($\mathrm{Prim}(3) = 30$) and $a = w_3(p_2) = 2e_2^{(3)}$. The idempotent $e_2^{(3)}$ satisfies $e_2^{(3)} \equiv 0 \pmod 2$, $e_2^{(3)} \equiv 1 \pmod 3$, $e_2^{(3)} \equiv 0 \pmod 5$; explicitly $e_2^{(3)} = 10 \bmod 30$, so $a = 20$. CRT components: $(a_1, a_2, a_3) = (0, 2, 0)$ in $\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/3\mathbb{Z} \times \mathbb{Z}/5\mathbb{Z}$. The classical Kronecker symbol gives $\left(\frac{20}{30}\right)_{\mathrm{Kron}} = 0$ (since $\gcd(20, 30) = 10 > 1$). The spectral Legendre symbol gives $\mathrm{SpecLeg}_3(20) = \chi_3(2) = \left(\frac{2}{3}\right) = -1$, recovering the local Legendre factor at $p_2 = 3$. This is the localisation in action.

Remark 6.4 ($p = 2$ convention and the X-label). The Kronecker value $\chi_2(a) = 0$ for even a is consistent with our convention $\mathrm{Label}_n(2) := X$ at $p_i = 2$: when we evaluate $\mathrm{SpecLeg}_n$ on the spectral weight $w_n(2) = 2e_1^{(n)}$, the index $k = 1$ (at $p_1 = 2$) is excluded (since $w_n(2)$ is even); the index $k = 2, \ldots, n$ are excluded (since $p_k \mid e_1^{(n)}$ for $k \geq 2$ by CRT orthogonality, hence $p_k \mid w_n(2)$). The empty-product convention gives $\mathrm{SpecLeg}_n(w_n(2)) = 1$, which we interpret via the X-label (a balanced / neutral value, not numerically $+1$ or -1). The classifier is therefore well-defined at every prime including 2, with the X-label acting as a neutral element whenever the classifier is extended multiplicatively to $\{B, C, X\}$-valued characters (as in the Kronecker-symbol extension of Remark 6.3).

Lemma 6.5 (Spectral Legendre reduction). *For every prime $p_i > 2$ and every depth $n \geq i$,*

$$\mathrm{SpecLeg}_n(w_n(p_i)) \;=\; \left(\frac{2}{p_i}\right). \tag{26}$$

Proof. Recall $w_n(p_i) = 2e_i^{(n)}$ with CRT-components determined by (15): $w_n(p_i) \equiv 2 \pmod{p_i}$ and $w_n(p_i) \equiv 0 \pmod{p_k}$ for all $k \leq n$ with $k \neq i$. Additionally $w_n(p_i)$ is even (divisible by 2), so if $i \neq 1$ then $p_1 = 2 \mid w_n(p_i)$ and the $k = 1$ index is excluded. The set of indices k with $p_k \nmid w_n(p_i)$ is therefore $\{i\}$: the index i is the unique one at which the weight has a nonzero CRT-component (namely $2 \bmod p_i$; since $p_i > 2$, $2 \not\equiv 0 \pmod{p_i}$). Hence

$$\mathrm{SpecLeg}_n(w_n(p_i)) \;=\; \chi_{p_i}(w_n(p_i) \bmod p_i) \;=\; \chi_{p_i}(2) \;=\; \left(\frac{2}{p_i}\right), \tag{27}$$

where the last equality uses $\chi_{p_i}(2) = \left(\frac{2}{p_i}\right)$ for $p_i > 2$ (by definition (24)). ☐

Alternative: direct CRT-local formulation.. The reader preferring to avoid a novel symbol entirely may take the following equivalent definition of the classifier, which uses only the single local component at p_i:

Definition 6.6 (CRT-local bipolar classifier). *For each prime $p_i > 2$ and each depth $n \geq i$, the* CRT-local classifier *is*

$$\mathrm{Label}_n(p_i) \;:=\; \begin{cases} B & \text{if } w_n(p_i) \bmod p_i \text{ is a nonzero quadratic residue modulo } p_i, \\ C & \text{if } w_n(p_i) \bmod p_i \text{ is a nonzero quadratic non-residue modulo } p_i; \end{cases} \tag{28}$$

$\mathrm{Label}_n(2) := X$.

Lemma 6.7 (Equivalence). *For every odd prime p_i and every depth $n \geq i$, Definition 6.6 gives $\mathrm{Label}_n(p_i) = B$ iff $\left(\frac{2}{p_i}\right) = +1$, and $\mathrm{Label}_n(p_i) = C$ iff $\left(\frac{2}{p_i}\right) = -1$. In particular, Definition 6.6 agrees with the $\mathrm{SpecLeg}$-based classifier of (26).*

Proof. $w_n(p_i) \bmod p_i = 2 \bmod p_i$ by the $k = i$ congruence of (15). Then $w_n(p_i) \bmod p_i$ is a quadratic residue modulo p_i iff $\left(\frac{2}{p_i}\right) = +1$, and a non-residue iff $\left(\frac{2}{p_i}\right) = -1$. ☐

Lemmas 6.5 and 6.7 together close the Jacobi-convention gap: either formulation is well-defined and recovers the Legendre-symbol classifier from the spectral weight. The remainder of the τ-derivation proceeds under whichever formulation the reader prefers.

6.5 Label convergence and the limiting classifier

Theorem 6.8 (Label convergence, III.T13)**.** *For every prime* p_i*, the sequence* $\mathrm{Label}_n(p_i)$ *for* $n \geq i$ *is constant:* $\mathrm{Label}_n(p_i) = \mathrm{Label}_i(p_i)$ *for all* $n \geq i$*. Consequently the limit*

$$\mathrm{Label}_\infty : \mathbb{P} \to \{B, C, X\}, \qquad \mathrm{Label}_\infty(p_i) \;:=\; \mathrm{Label}_i(p_i) \tag{29}$$

is well defined, and for odd p*,* $\mathrm{Label}_\infty(p) = B$ *if* $p \equiv \pm 1 \pmod 8$*,* $\mathrm{Label}_\infty(p) = C$ *if* $p \equiv \pm 3 \pmod 8$*;* $\mathrm{Label}_\infty(2) = X$*.*

Proof. By Lemma 6.5 (equivalently, Lemma 6.7), for each $n \geq i$ the classifier $\mathrm{Label}_n(p_i)$ reads off the single value $\left(\frac{2}{p_i}\right)$, which is a function of p_i alone and does not depend on n. Hence $\mathrm{Label}_n(p_i) = \mathrm{Label}_i(p_i)$ for all $n \geq i$, so the limit is well defined and equals the value at the first admissible depth. The mod-8 form follows from Remark 5.2. □

This completes the τ-framework derivation: the internal bipolar classifier is the Legendre-symbol function $\left(\frac{2}{\cdot}\right)$, reached structurally by CRT decomposition of the primorial ring, split-complex extension, and the spectral-weight-to-localised-Legendre reduction (Theorem 6.5).

7. THE ISOMORPHISM THEOREM

Theorem 7.1 (Isomorphism)**.** *The orthodox classifier* Pol *(Definition 5.1) and the τ-framework classifier* Label_∞ *(Theorem 6.8) agree pointwise:*

$$\mathrm{Pol}(p) \;=\; \mathrm{Label}_\infty(p) \qquad \textit{for all primes } p. \tag{30}$$

Consequently $\mathbb{P}_B^{\mathrm{Pol}} = \mathbb{P}_B^{\mathrm{Label}_\infty}$ *and* $\mathbb{P}_C^{\mathrm{Pol}} = \mathbb{P}_C^{\mathrm{Label}_\infty}$*, and the orthodox Prime Polarity Theorem (Theorem 5.3) transfers verbatim to the τ-framework classifier.*

Proof. For $p = 2$, both Pol and Label_∞ return X by direct definition. For odd prime p, both return $\left(\frac{2}{p}\right)$ translated to $\{B, C\}$ by the same convention ($+1 \mapsto B$, $-1 \mapsto C$): Pol by Definition 5.1, and Label_∞ by Theorem 6.8 (via Lemma 6.5). □

7.1 Why the isomorphism matters

The Isomorphism Theorem is not a re-proof of Gauss. Its value lies in the *bridge* it establishes between two frameworks that develop prime polarity through entirely different structural routes.

From τ to classical.. The τ-side derivation begins with kernel axioms, extracts a CRT decomposition of the primorial ring (III.T10), extends by split-complex scalars to produce the spectral weight $w_n(p_i)$ (III.D23), and reads off the bipolar label via the SpecLeg reduction (Theorem 6.5). The isomorphism says that the classifier so obtained is the *same* as the Legendre-symbol classifier that Gauss identified in 1801. For the τ-framework this is a consistency check of the highest order: the structural derivation, not being calibrated externally, could in principle have produced a different classifier. It did not.

From classical to τ.. The converse direction is what downstream τ-framework constructions require. The classical Legendre classification, while correct, carries no information about how the primes participate in the spectral algebra of the lemniscate boundary. The τ-derivation attaches each prime to its CRT idempotent $e_i^{(n)}$, hence to a specific element of the split-complex boundary ring $H_\tau \otimes \mathbb{Z}/\mathrm{Prim}(n)\mathbb{Z}$, hence to a specific character of the boundary-character algebra $\mathrm{Char}(\mathbb{L})$ used in Book III Chapters 19 and following [12]. The Isomorphism Theorem says that this attachment is consistent with the classical classification: when one asks “which character does prime p correspond to?”, the answer is $\left(\frac{2}{p}\right)$, and the choice of split-complex $\mathrm{j}^2 = +1$ rather than elliptic $i^2 = -1$ affects the *lobe structure* downstream but not the per-prime label.

Downstream use in the program.. This bridge is what allows the companion papers on black-hole stability [7] and arithmetic quantum gravity [6] to write L-function factorisations and spectral decompositions that mention the prime classes $\mathbb{P}_B, \mathbb{P}_C$ as operational objects. The present theorem guarantees that every such mention can be read either as a statement about Legendre residues (if the reader prefers a classical frame) or as a statement about CRT-local spectral characters (if the context of Book III’s spectral algebra is invoked). The two readings are mathematically identical and interchangeable.

Connection to the Riemann hypothesis programme.. Book III's Critical Line Theorem (III.T19, conditional on the determinant representation O3; see [6]) derives the critical-line property of τ-L-functions from self-adjointness of a spectral operator on the boundary character space. The input to that construction is the bipolar decomposition of the boundary algebra, which uses the classifier established here. The Prime Polarity Theorem is therefore not only a foundational result in its own right but also a load-bearing ingredient for the Riemann hypothesis programme in the τ-framework.

Corollary 7.2 (Unified theorem statement)**.** *Let* $\mathrm{Pol} \equiv \mathrm{Label}_\infty$ *denote the common classifier. Then:*

(1) Pol *is computable in* $O(\log p)$ *arithmetic operations.*

(2) $\mathbb{P} = \mathbb{P}_B \sqcup \mathbb{P}_C \sqcup \{2\}$ *as a disjoint partition.*

(3) $|\mathbb{P}_B| = |\mathbb{P}_C| = \aleph_0$.

(4) Each of $\mathbb{P}_B, \mathbb{P}_C$ *has natural density* $\frac{1}{2}$ *in the primes.*

The four claims are theorems of ZFC. Parts (1)–(2) are provable in PA; Part (3) (infinitude) is provable in RCA_0 *via the elementary Selberg–Erdős proof of Dirichlet; Part (4) (density) requires either analytic Dirichlet density (L-functions) or the Siegel–Walfisz form of PNT for arithmetic progressions. See Appendix A for the reverse-mathematical locator.*

The τ-framework side. *Under the spectral Legendre convention (§6.4), all four claims are also theorems of Category τ derived rigorously via the CRT-local reduction (Theorem 6.5 and the equivalent Lemma 6.7). Part (4) (density) combines the two Dirichlet residue classes of each polarity label (*$\{1,7\} \pmod 8$ *for B;* $\{3,5\} \pmod 8$ *for C) into density* $2 \cdot \frac{1}{\varphi(8)} = \frac{1}{2}$ *per class.*

8. DISCUSSION AND OPEN QUESTIONS

8.1 What we have, and what remains open

The paper delivers a fully rigorous Prime Polarity Theorem in both orthodox and τ-framework form, with a pointwise-equal classifier. The rigorous classifier is the Legendre symbol $(\frac{2}{\cdot})$; the growth-rate comparison of Book I Chapter 27 survives as *motivation*, not as the derivation.

The growth-rate story remains structurally coherent: the idempotent decomposition $H_\tau \otimes \mathbb{Q} \cong \mathbb{Q}e_+ \oplus \mathbb{Q}e_-$ is the split-complex algebraic counterpart of the two growth regimes, and the spectral weight $w_n(p_i) = 2e_i^{(n)}$ is what one obtains when one asks "how does the CRT idempotent project onto each lobe?" The sharp classifier comes from the Legendre symbol, but the *reason* the Legendre symbol is the right classifier is tied to the e_+/e_- decomposition.

The intrinsic-growth-rate question.. A more ambitious programme would be to recover the Legendre classification *directly* from intrinsic growth-rate dynamics of τ-Idx, without passing through the CRT idempotents. Book I Chapter 27's "Stabilization Mechanism" section sketches such a procedure via the scaled-gap primorial-ladder argmax:

$$\Delta\left(\frac{M_k}{p^p}\right) \cdot p^p \overset{?}{<} \Delta\left(\frac{M_k}{p}\right) \cdot p, \tag{31}$$

where M_k is the k-th primorial and Δ the prime-gap function. The conjecture would be that the eventual winner of (31) as $k \to \infty$ matches the Legendre classification. We pose this as Open Question 1 (§8).

Our current view is that an affirmative answer is plausible but requires deep prime-gap analysis. The needed ingredients are refinements of the Maynard–Tao small-gaps results [18, ch. 9] on the C-dominance side (to show that small gaps at M_k/p^p occur infinitely often, along suitable subsequences) and Rankin-type large-gap results on the B-dominance side (to show that comparable or larger gaps occur at M_k/p). The match with the Legendre classification would then be a prime-gap statement about residue structure modulo 8 that we believe is accessible but have not undertaken here. We stress that an intrinsic-growth-rate proof would *not* affect the correctness of the Legendre classification; it would provide an additional structural explanation.

8.2 Open questions

(OQ1) **Intrinsic growth-rate classifier.** An intrinsic-growth-rate classifier at the k-th primorial stage is defined via the scaled-gap inequality

$$\Delta(\lfloor M_k/p^p \rfloor) \cdot p^p \overset{?}{<} \Delta(M_k/p) \cdot p, \tag{32}$$

where $M_k = \text{Prim}(k)$ is the k-th primorial, $\Delta(m) := m - \text{prevPrime}(m)$ is the prime-gap function at the integer m, and $\lfloor \cdot \rfloor$ is applied because $\text{Prim}(k)$ is squarefree so M_k/p^p is non-integer for any $p \geq 2, k \geq 1$. (Note $M_k/p \in \mathbb{N}$ for k at least the index of p.) The conjecture below is that the winner of (32) stabilises as $k \to \infty$ and that the eventual winner gives the B/C polarity of p:
Under the sign convention that B *wins at stage* k when $\Delta(\lfloor M_k/p^p \rfloor) \cdot p^p > \Delta(M_k/p) \cdot p$ and C *wins* otherwise, define for $k \geq k_0(p) := \min\{k : \lfloor M_k/p^p \rfloor \geq 3\}$

$$s_k(p) \;:=\; \text{sign}\Big(\Delta(\lfloor M_k/p^p \rfloor) \cdot p^p \;-\; \Delta(M_k/p) \cdot p\Big) \in \{+1, -1, 0\}. \tag{33}$$

Conjecture 8.1 (Scaled-gap classifier matches Legendre)**.** *For every odd prime p:*

(a) if $\text{Pol}(p) = B$ *(equivalently* $p \equiv \pm 1 \pmod 8$*), the sequence* $s_k(p)$ *stabilises eventually to* $+1$*;*

(b) if $\text{Pol}(p) = C$ *(equivalently* $p \equiv \pm 3 \pmod 8$*),* $s_k(p) = -1$ *along a subsequence of positive asymptotic density in* k*.*

The asymmetric form of Conjecture 8.1 reflects the density asymmetry of the scaled-gap inequality: the B-channel scales by p^p against the C-channel's p, so the B-channel dominates *typically*; the C-dominance condition becomes a statement about unusual prime-gap coincidences occurring infinitely often (and, per (b), at positive density) along specific residue structures modulo 8.

Numerical evidence (Table 2). For accessible $k \leq 12$, we computed $s_k(p)$ for $p \in \{3, 5, 7, 11\}$ (the first four odd primes for which $k_0(p)$ is reached before $k = 12$). The results are consistent with Conjecture 8.1(a) (every sampled B-prime has $s_k = +1$ throughout the range). For Conjecture 8.1(b), the C-dominance condition is observed only once in the sampled range (at $p = 3, k = 11$), consistent with the density asymmetry but inconclusive at accessible k. A larger-scale computation verifying (b) is an attractive computational open problem.

Table 2. Computed sign $s_k(p)$ from (33) for selected (p, k). "·" indicates $k < k_0(p)$. Consistent with Conjecture 8.1(a): B-primes $\{7, \ldots\}$ show $s_k = +1$ throughout. Consistent with Conjecture 8.1(b) only sparsely: the C-dominance value $s_k = -1$ is observed at $(p, k) = (3, 11)$, a single positive sample among 29 computed cells. Larger-scale verification of (b) is open.

p (Pol)	k = 4	5	6	7	8	9	10	11	12
3 (C)	+	+	+	+	+	+	+	−	+
5 (C)	·	·	+	+	+	+	+	+	+
7 (B)	·	·	·	·	+	+	+	+	+
11 (C)	·	·	·	·	·	·	·	·	+

Proving Conjecture 8.1(b) would require deep prime-gap analysis: Maynard–Tao small-gaps-in-AP results [18, Ch. 9] on the C-dominance side, Rankin-type large-gap results on the B-dominance side, and a match with the mod-8 structure of the Legendre classification. None of these ingredients is individually out of reach; a unified proof has not been attempted.

(OQ2) **Density of $\mathbb{P}_B^{\text{growth}}, \mathbb{P}_C^{\text{growth}}$.** Assuming (OQ1) has an affirmative answer, the densities of the intrinsic $\mathbb{P}_B, \mathbb{P}_C$ classes would match those of Pol (density $\frac{1}{2}$ each). Could this be proved without going through the Legendre symbol, using only prime-gap analysis?

(OQ3) **Higher-rank generalisation.** The classifier Pol has codomain $\{B, C, X\}$; the "mixed" class $\mathbb{P}_X$ consists of the single prime $\{2\}$. Is there a generalisation where the codomain grows with further structural layers of Category τ? The 4+1-sector decomposition of III.D13 [12] is a candidate target.

(OQ4) **Full Sato–Tate for τ-L-functions.** In our companion paper on arithmetic quantum gravity, we flagged obligation O4: derive the $\sin^2\theta$ Sato–Tate distribution for Hecke eigenvalue angles from the τ-framework. The present Legendre classifier gives density $\frac{1}{2}$ (Sato–Tate's zeroth moment); the question is whether the full distribution is derivable from the split-complex lobe structure.

(OQ5) **Lean certification.** The Lean plan of §B.2 remains to be executed.

9. CONCLUSION

We have proved the Prime Polarity Theorem in orthodox ZFC/PA form using Euler's criterion and Dirichlet's theorem (Theorem 5.3), re-derived the same classifier inside Category τ from the CRT idempotents on the primorial ring extended by the split-complex boundary ring (Theorem 6.8), and shown that the two classifiers coincide pointwise (Theorem 7.1). The proof rests entirely on classical mathematics; the τ-contribution is the *structural derivation* of the classifier from first principles of the categorical framework, rather than its *postulation*.

Along the way we verified that neither the universal-existential growth-rate criterion (Lemma 2.1) nor the bound-dependent spectral signature (Lemma 3.1) is sufficient to realise the theorem: the first is vacuous, the second has asymptotically empty C-class. The Legendre-symbol classifier is the unique surviving candidate among the natural alternatives, and the Isomorphism Theorem shows that the structural derivation recovers precisely this classifier.

The Prime Polarity Theorem is thus a theorem of classical number theory with a clean τ-framework derivation and a tractable Lean-formalisation path. The partition structure makes the downstream chain of the Panta Rhei program coherent, and the Isomorphism Theorem establishes the bridge to analytic number theory that is needed for the L-function and spectral-algebra constructions in Book III and the companion papers [7, 6].

ACKNOWLEDGEMENTS

We thank the Lean 4 and `mathlib` communities [19, 22], whose computable prime-polarity implementations enabled the empirical diagnostic in §4 and whose analytic- and algebraic-number-theory primitives underwrite the formalisation pathway outlined in Appendix B.2. The reverse-mathematical locator in Appendix A follows the standard framework of Simpson [21], with the elementary proof of the Prime Number Theorem due to Avigad, Donnelly, Gray, and Raff [2]. The integer-multiplication complexity bound used in Proposition 5.5 is that of Harvey and van der Hoeven [16].

Conjecture 8.1 and its numerical-evidence table (Table 2) would not be formulable without the prime-gap computational machinery of the `sympy` library, which we used for the direct computation of $s_k(p)$ at small (p, k). We thank several discussions surrounding the companion papers [7, 6] for shaping the presentation of the Isomorphism Theorem's downstream use.

A. REVERSE-MATHEMATICAL LOCATOR

A pedantically careful reader will want to know which fragment of arithmetic each part of Theorem 5.3 actually needs. We provide a compact locator. The general framework and results are from Simpson's standard reference on reverse mathematics [21].

Summary. The *classification* (parts 1–2) is elementary and low-complexity, at home in PA. The *infinitude* (part 3) is mid-strength, provable in RCA_0 via the Selberg–Erdős route. The *density* (part 4) sits at the strength of PNT for arithmetic progressions, accommodated in WKL_0. None of the parts requires full ZFC; the paper's statement as a ZFC theorem is for convenience, not necessity. Exact localisations of parts (3) and (4) within reverse mathematics remain open and may be of independent interest.

B. LEAN 4 FORMALISATION STATUS AND PLAN

B.1 What the Lean library actually establishes

We have inspected the TauLib modules `TauLib/BookI/Polarity/{Polarity, Spectral}.lean` at the commit-of-record. The Lean content falls into three categories.

Category A — formally proved theorems..

- `b_channel_unbounded` (`Polarity.lean:59--61`): $\forall a \geq 2, \forall \text{target}, \exists B.\ a^B > \text{target}$.
- `c_channel_unbounded` (`Polarity.lean:65--66`): $\forall a \geq 2, \forall \text{target}, \exists C.\ a \uparrow\uparrow C > \text{target}$.
- `growth_rate_separation` (`Spectral.lean:77--78`): $\forall a \geq 2, \forall B, \exists C.\ a \uparrow\uparrow C > a^B$.
- `b_beats_c` (`Polarity.lean:239--244`): $\forall a \geq 2, \forall C, \exists B.\ a^B > a \uparrow\uparrow C$.
- `exp_strict_mono` (`Polarity.lean:70--75`): strict monotonicity of exponentiation for $a \geq 2$.
- `pow_ge_succ` (`Polarity.lean:45--52`): $a^n \geq n + 1$ for $a \geq 2$.

Table 3. Axiomatic strength of the four parts of Theorem 5.3. PA: Peano arithmetic. RCA_0: recursive comprehension. WKL_0: weak König's lemma over RCA_0. ACA_0: arithmetic comprehension. Upper bounds are stated without claiming exhaustiveness; tighter localisations may be possible and are not pursued here.

Part	Upper bound	Remarks
(1) Partition $\mathbb{P} = \mathbb{P}_B \sqcup \mathbb{P}_C \sqcup \{2\}$	PA (in fact $I\Sigma_1$)	Follows from computability of Pol (Prop. 5.5) and the mod-8 case analysis of Remark 5.2.
(2) Dichotomy for odd primes	PA	Euler's criterion and the second supplement of QR are provable in $I\Sigma_1$; the mod-8 case analysis is primitive recursive.
(3) Infinitude of $\mathbb{P}_B$ and $\mathbb{P}_C$	RCA_0 (elementary); WKL_0 (analytic)	Infinitude of primes in a fixed residue class mod 8 is classical (Dirichlet, 1837). The Selberg–Erdős elementary route formalises in RCA_0. The original analytic Dirichlet route uses L-functions and is accommodated in WKL_0 with the standard second-order coding. See Simpson [21] §IV for the general framework.
(4) Natural density $\frac{1}{2}$ each	WKL_0 (sufficient)	Equivalent to the Prime Number Theorem for arithmetic progressions (Dirichlet / de la Vallée Poussin / Siegel–Walfisz). The base PNT has a formally verified proof in elementary (RCA_0-style) methods [2]; the arithmetic-progression refinement is accommodated in WKL_0 along the same lines. The Dirichlet-density refinement (weaker claim) holds already in RCA_0.

Category B — computable definitions and decision procedures..

- `spectral_sig`, `b_max`, `c_max` (`Spectral.lean:35--58`).
- `pol_at`, `polarity_map`, `polarity_chi` (`Polarity.lean:119--138`).
- `chi_extend` (multiplicative extension; `Polarity.lean:145--161`).
- `b_class_witness`, `c_class_witness`, `count_b_dominant`, `count_c_dominant`, `partition_check` (`Polarity.lean:173--221`).

Category C — claimed but not present as theorems..

- *Dichotomy* "every prime is B-dominant or C-dominant": no Lean `theorem` of this content.
- *B-class infinite*: no Lean `theorem`.
- *C-class infinite*: no Lean `theorem`.
- *Hyperfactorization uniqueness* (I.T04): proved rigorously in the bundle-companion Hyperfactorization paper [8] (Hinge 1), with full details also in Book I Chapter 24 manuscript. It is not yet formalised as a `theorem` in TauLib's `Coordinates/Hyperfact.lean` (which contains only computable checks `hyperfact_check`, `encoding_check`); formalisation is a follow-up target.

B.2 Proof-chain sketch

We give a structural sketch of how each part of the main theorem traces to existing `mathlib` primitives, so that a reader can judge the plausibility of a Lean 4 certification without requiring commit-level detail. We work with the assumption that `mathlib` identifiers continue to name the objects described; exact names may differ at implementation time.

Step 1: Definition of Pol**..** A decidable function `Pol : Nat` → `Polarity` (with `Polarity` an inductive type $\{B, C, X\}$) computes $\mathrm{Pol}(p)$ by mod-8 case analysis of Remark 5.2. The required primitive is `ZMod`'s computable mod-8 arithmetic (`Mathlib/Data/ZMod`) and a two-line inductive type in `TauLib`.

Step 2: Agreement of Pol **with Legendre symbol..** The equality $\mathrm{Pol}(p) = \left(\frac{2}{p}\right)$ (for odd p) reduces to `legendreSym.at_two` (`Mathlib/NumberTheory/LegendreSymbol/QuadraticReciprocity`) giving $\mathrm{legendreSym}\, p\, 2 = \chi_8(p)$, combined with the case-analysis Step 1. This pipeline is the second supplement of quadratic reciprocity already packaged in `mathlib`.

Step 3: Dichotomy for odd primes.. `theorem Pol_dichotomy`: for every odd prime p, $\mathrm{Pol}(p) \in \{B, C\}$. Direct consequence of Step 1 and decidability of p mod 8.

Step 4: Infinitude of $\mathbb{P}_B, \mathbb{P}_C$.. Relies on `mathlib`'s Dirichlet-in-AP result (`Nat.infinite_setOf_prime_and_eq_mod` in `Mathlib/NumberTheory/LSeries/PrimesInAP`), applied to the four residue classes $\{1, 3, 5, 7\} \pmod 8$ and summed.

Step 5: Density.. The density claim relies on `mathlib`'s analytic-number-theory suite (the `DirichletL` family under `Mathlib/NumberTheory/LSeries/`). At present, `mathlib` contains the ingredients (Dirichlet L-functions, non-vanishing at $s = 1$) but not the packaged statement "density $\frac{1}{\varphi(q)}$ per residue class". A first Lean pass may cite this density as an external fact; a later pass formalises it once `mathlib` matures its packaging.

Step 6: Spectral Legendre reduction.. Theorem 6.5 admits two Lean routes:

- *Route A* (SpecLeg direct): define `SpecLeg_n : ZMod (Prim n) → Int` using `ZMod.prodEquivPi` (`Mathlib/Data/ZMod/QuotientRing`) for the CRT decomposition and pointwise `legendreSym` at each local component.
- *Route B* (CRT-local, Definition 6.6): use `IsSquare` on `ZMod p` via `legendreSym.eq_one_iff`. Route B is simpler and *independent of Hyperfactorization* (I.T04), making it the preferred Lean route for the τ-side proof. The independence holds because Route B's entire pipeline (CRT idempotent $e_i^{(n)}$, spectral weight $w_n(p_i) = 2e_i^{(n)}$ via the reduced trace, $\chi_{p_i}(w_n(p_i) \bmod p_i)$) lives at the CRT-idempotent level in $H_\tau \otimes \mathbb{Z}/\mathrm{Prim}(n)\mathbb{Z}$, without ever invoking the hyperfactorization decomposition $X = T(A, B, C) \cdot D$ of arbitrary integers. The coordinate functions $\mathrm{coord}_A, \mathrm{coord}_B, \mathrm{coord}_C$ enter this paper only in the Lemma 2.1 and Lemma 3.1 diagnostics, not in the main theorem's Lean pipeline.

Under either route, the reduction $\mathrm{SpecLeg}_n(w_n(p_i)) = (\frac{2}{p_i})$ is immediate from the CRT-orthogonality congruences $w_n(p_i) \equiv 2 \pmod{p_i}$ and $w_n(p_i) \equiv 0 \pmod{p_k}$ for $k \neq i$.

Step 7: Label convergence.. Theorem 6.8 follows from Step 6: $\mathrm{Label}_n(p_i) = \chi_{p_i}(2) = (\frac{2}{p_i})$ is a function of p_i alone, so $\mathrm{Label}_n(p_i)$ is constant in n for $n \geq i$. The Lean formalisation is a one-line "`Label_n_eq_Label_i`" lemma.

Step 8: Isomorphism Theorem.. Theorem 7.1 combines Steps 2 and 6: both Pol and Label_∞ evaluate to $\chi_8(p)$ on odd primes (and to X at $p = 2$), hence coincide pointwise. Under Route B of Step 6 the combination is a direct `Decidable`-case analysis and requires no further imports.

Summary.. The Lean-formalisation gap between this paper and a machine-checked proof is a single sprint's worth of work once `mathlib`'s Dirichlet-in-AP infrastructure (`PrimesInAP.lean`) is in the imported set, with the density claim (Step 5) being the only part possibly deferred. The τ-side (Steps 6–8) is approximately 200 Lean lines assuming Route B; an alternative route via the Hyperfactorization Theorem (I.T04; formalised in the companion paper) adds depth but is not strictly required for Step 6. The line-count estimates are $\approx$ 300–400 Lean lines for the orthodox side (Steps 1–5) and $\approx$ 200 Lean lines for the τ-side (Steps 6–8, Route B).

REFERENCES

[1] Tom M. Apostol. *Introduction to Analytic Number Theory*. Undergraduate Texts in Mathematics. Springer, New York, 1976.

[2] Jeremy Avigad, Kevin Donnelly, David Gray, and Paul Raff. A formally verified proof of the prime number theorem. *ACM Transactions on Computational Logic*, 9(1):Article 2, 2007.

[3] Nikolai Chebotarëv. Die bestimmung der dichtigkeit einer menge von primzahlen, welche zu einer gegebenen substitutionsklasse gehören. *Mathematische Annalen*, 95:191–228, 1926.

[4] Henri Cohen. *A Course in Computational Algebraic Number Theory*, volume 138 of *Graduate Texts in Mathematics*. Springer, Berlin, 1993.

[5] Peter Gustav Lejeune Dirichlet. Beweis des satzes, dass jede unbegrenzte arithmetische progression .... *Abhandlungen der Königlich Preußischen Akademie der Wissenschaften*, pages 45–81, 1837.

[6] Thorsten Fuchs and Anna-Sophie Fuchs. Arithmetic quantum gravity without singularities: A categorical reading of the conformal primon gas, 2026. Panta Rhei Research, `https://panta-rhei.site/papers/primon-categorical`.

[7] Thorsten Fuchs and Anna-Sophie Fuchs. Black hole stability without extra dimensions: A categorical reinterpretation of the G_2-manifold remnant, 2026. Panta Rhei Research, `https://panta-rhei.site/papers/bh-categorical`.

[8] Thorsten Fuchs and Anna-Sophie Fuchs. The hyperfactorization theorem: Unique tower-atom decomposition in ZFC and category τ, 2026. Panta Rhei Research, `https://panta-rhei.site/papers/hyperfactorization`. Hinge 1 of the Panta Rhei bundle.

[9] Thorsten Fuchs and Anna-Sophie Fuchs. The master constant ι_τ: A structural derivation of $\iota_\tau = 2_\tau/(\pi_\tau + e_\tau)$, 2026. Panta Rhei Research, `https://panta-rhei.site/papers/iota-tau`. Hinge 3 of the Panta Rhei bundle.

[10] Thorsten Fuchs and Anna-Sophie Fuchs. *Panta Rhei, Book I: Categorical Foundations*. Panta Rhei Research, 2nd edition, 2026. ISBN 979-8-3462-1987-6.

[11] Thorsten Fuchs and Anna-Sophie Fuchs. *Panta Rhei, Book II: Categorical Holomorphy*. Panta Rhei Research, 2nd edition, 2026. ISBN 979-8-3462-1988-3.

[12] Thorsten Fuchs and Anna-Sophie Fuchs. *Panta Rhei, Book III: Categorical Spectrum*. Panta Rhei Research, 2nd edition, 2026. ISBN 979-8-3462-1989-0.

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

[14] Carl Friedrich Gauss. *Disquisitiones Arithmeticae*. 1801. English translation by A. A. Clarke, Springer, 1986.

[15] G. H. Hardy and E. M. Wright. *An Introduction to the Theory of Numbers*. Oxford University Press, Oxford, 6 edition, 2008. Revised by D. R. Heath-Brown and J. H. Silverman.

[16] David Harvey and Joris van der Hoeven. Integer multiplication in time $o(n \log n)$. *Annals of Mathematics*, 193(2):563–617, 2021.

[17] Kenneth Ireland and Michael Rosen. *A Classical Introduction to Modern Number Theory*, volume 84 of *Graduate Texts in Mathematics*. Springer, New York, 2 edition, 1990.

[18] Henryk Iwaniec and Emmanuel Kowalski. *Analytic Number Theory*, volume 53 of *American Mathematical Society Colloquium Publications*. American Mathematical Society, Providence, RI, 2004.

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

[20] Jean-Pierre Serre. *A Course in Arithmetic*, volume 7 of *Graduate Texts in Mathematics*. Springer, New York, 1973.

[21] Stephen G. Simpson. *Subsystems of Second Order Arithmetic*. Perspectives in Logic. Cambridge University Press and ASL, New York and La Jolla, 2 edition, 2009.

[22] The mathlib Community. The Lean mathematical library. Software and documentation, 2024. `https://leanprover-community.github.io/`.