

Structural Construction of the Uniqueness Statement for the Riemann Hypothesis A Chain Map between the Additive and Multiplicative Sides, a Well-Defined Obstruction Functional, and a Spectral Uniqueness Theorem

Yang Liu¹

1620872416@qq.com

¹Independent researcher

Article

Keywords: Riemann hypothesis, chain map, rank gap, spectral theorem, explicit formula, Li criterion, Weil positivity, Chern character, fundamental theorem of arithmetic

Posted Date: October 9th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11291222/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.
[Read Full License](#)

Additional Declarations: No competing interests reported.

Structural Construction of the Uniqueness Statement for the Riemann Hypothesis

A Chain Map between the Additive and Multiplicative Sides, a Well-Defined Obstruction Functional, and a Spectral Uniqueness Theorem

Yang Liu^{1*}

¹Independent researcher, Chengdu, China

*Correspondence: 1620872416@qq.com

Abstract

The Riemann zeta function admits two presentations whose equality is pure arithmetic: the additive Dirichlet series $\sum_{n \geq 1} n^{-s}$ and the multiplicative Euler product $\prod_p (1 - p^{-s})^{-1}$ coincide, for $\text{Re } s > 1$, by the fundamental theorem of arithmetic alone. Yet the free structures that carry the two presentations have different rank — rank 1 for the additive side, $\text{rank } |\mathbb{P}| = \infty$ for the multiplicative side — so the numerical identity of the two presentations is not accompanied by any structural identity. This rank gap is the precise locus at which the Riemann Hypothesis (RH) acquires structural, as opposed to merely numerical, content. We carry the corresponding structural programme to completion and isolate a single remaining item.

Six constructions are established. (i) **A chain map.** Using the fundamental theorem of arithmetic we construct the explicit map $\varphi(n) = \sum_p v_p(n)[p]$ from the multiplicative side into the free abelian group on the primes, and prove that it is a group homomorphism and an injection but not a surjection. Whence the **rank-gap theorem**: no chain map between the two sides induces an isomorphism on homology, so the requirement that the two sides be *homotopy equivalent* is not merely unproven but impossible. (ii) **A well-defined obstruction functional.** The obstruction δ_ζ , previously written as the pointwise series $\sum_p (\log p) \theta_p / \pi$, which diverges on the critical line, is re-defined in the cumulative form $\delta_\zeta(x) = \psi(x) - x + \log 2\pi + 1/2 \log(1 - x^{-2}) = -\sum_p x^\rho / \rho$. Its two sides — primes and zeros — agree numerically, with relative residual decaying from 1.64×10^{-2} at $x = 10^2$ to 1.06×10^{-4} at $x = 10^5$. (iii) **A spectral uniqueness theorem.** If a self-adjoint operator H has the non-trivial zero ordinates as its spectrum, then H is unique up to unitary equivalence; a Jordan-block counterexample shows that self-adjointness is load-bearing, and the eigenvalue counting function is locked to the Riemann–von Mangoldt density, so the operator carries no continuous free parameter. (iv) **A topological bridge.** The Chern character $\text{ch} = (\text{rank}, c_1)$ maps both sides into a common cohomology ring and satisfies a discrete Gauss–Bonnet identity, providing the topological dual of the analytic (explicit-formula) bridge. (v) **A structural baseline.** The critical line is not assumed. The value $1/2$ is returned independently by the reflection symmetry of the completed zeta function (analytic route) and by the self-normalisation exponent of the quadratic Gauss sum, $|G(p)| = \sqrt{p}$ (arithmetic route, using no analytic input); the two coincide, so $\text{Re } s = 1/2$ is a structural output rather than a convention. Moreover the discrete reciprocity character lifts to the continuous local Euler phase only at this baseline — the *seam lift* — which closes the apparent type mismatch between the discrete symbol ± 1 and the continuous angle. (vi) **A dynamical route.** The critical line is characterised dynamically as well as geometrically: the de Bruijn–Newman backwards heat flow $H_t(z) = \int_0^\infty e^{-(tu^2)} \Phi(u) \cos(zu) du$ (with $H_0 = 1/8 \xi(1/2 + iz/2)$) has all real zeros exactly for $t \geq \Lambda$, with $\text{RH} \Leftrightarrow \Lambda = 0$; and its zero dynamics is the gradient flow of the logarithmic Hamiltonian $H(t) = \sum_{j \neq k} (j - k) \log(1/|x_j - x_k|)$, whose energy $E(t) = \sum_{j \neq k} |x_j - x_k|^{-2}$ is monotone non-increasing, so that least dissipation forces the configuration to an equal-spacing equilibrium (Rodgers–Tao; Polymath 15). The route is read in reverse: the flow necessarily

possesses a baseline, but it moves in the real variable and cannot select the vertical line — that baseline is the value $1/2$ already fixed, independently, by the spatial argument of §7.7, so the dynamics meets it rather than deriving it. This localises the open item to a **single real number**, bracketed unconditionally by $0 \leq \Lambda \leq 0.22$, and reduces it to the assertion that the arithmetic datum already sits on the baseline (§6.7). The two quantities this reading rests on are formalised in §6.8: the **dissipation work** $W = \int \text{Edt} = 1/4\Delta H$, and the **reciprocal efficiency** $\eta = E_{\text{lat}}/E_{\text{act}} \in (0,1]$, maximal ($\eta = 1$) exactly at the density-matched equal-spacing ground state.

The remaining item — the Hilbert–Pólya existence statement — is shown not to be a notion without content: it is equivalent to the Li positivity criterion $\lambda_n \geq 0$ for all n and to Weil positivity, and the Li partial sums are numerically positive for $n \leq 20$. It is moreover **localised**, not left as an unlocated gap. In two directly comparable structures — function fields (Weil 1948) and compact hyperbolic surfaces (Selberg 1956) — the analogue of RH is a *theorem*, because there the relevant operator is self-adjoint for free, by a positive-definite pairing (an ample polarisation, a Riemannian metric). The arithmetic case has the same schematic form with that positivity missing, and supplying it is exactly RH (§6.6). It is localised a second time, dynamically, to the single scalar $\Lambda = 0$ of de Bruijn–Newman, bracketed unconditionally by $0 \leq \Lambda \leq 0.22$ (§6.7). The discrete zero positions on those structures are computed and compared: for 123 elliptic curves over $\mathbb{F}_p$ every Frobenius eigenvalue satisfies $|\alpha| = \sqrt{p}$ and every zero lies on $\text{Re } s = 1/2$ (worst deviation 1.1×10^{-16}), matching the line and the mechanism, while the density laws differ (Weyl r^2 against Riemann–von Mangoldt). Every item recorded as a gap in an earlier treatment is accounted for: eleven of twelve are closed by a named theorem or proposition — including the items earlier described only as "limitations", which are upgraded to non-decidability theorems — and the twelfth is ES itself, now localised (§10.11). The paper therefore reduces RH, along this route, to one sharply stated existence problem with a computable positivity form, a named localisation, and a two-sided bracket, and reports the exact boundary of what the construction does and does not settle.

Keywords: *Riemann hypothesis; chain map; rank gap; spectral theorem; explicit formula; Li criterion; Weil positivity; Chern character; fundamental theorem of arithmetic*

1 Introduction

1.1 Background

The Riemann Hypothesis (RH) asserts that every non-trivial zero of the Riemann zeta function $\zeta(s)$ lies on the critical line $\text{Re } s = 1/2$. Since Riemann's 1859 memoir [1], in which the function, its analytic continuation, its functional equation and its connection with the distribution of primes were introduced, the conjecture has generated an enormous body of equivalent formulations, partial results and numerical verification. von Mangoldt [2] established the explicit formula that expresses the prime-counting function in terms of the zeros; Hardy and Littlewood proved that infinitely many zeros lie on the critical line; Selberg [3] and later Connes [4] developed trace-formula interpretations; Montgomery [5] identified the statistical behaviour of the zero spacings with the eigenvalue statistics of random Hermitian matrices; and Platt and Trudgian [6] have rigorously verified the hypothesis for the first 3×10^{12} zeros. No proof is known, and the problem stands among the Clay Millennium Prize Problems [7].

The present paper is not concerned with sharpening partial results, with improving zero-free regions, or with extending numerical verification. It is concerned with a structural observation that is elementary — it uses nothing beyond the fundamental theorem of arithmetic — but that, as we shall argue, is decisive for the way the problem should be posed. The observation is that the zeta function has two presentations,

$$\zeta(s) = \sum_{n \geq 1} n^{-s} \quad \text{and} \quad \zeta(s) = \prod_{p \in \mathbb{P}} (1 - p^{-s})^{-1},$$

both valid in the half-plane $\text{Re } s > 1$, and that their equality — first in the common region of convergence, and then everywhere by analytic continuation [8, 9] — rests *only* on the unique factorisation of integers into primes [10]. This is a statement about numbers, not about analysis. It is therefore natural to ask what the two presentations *are*, structurally, and whether the arithmetic fact that they describe one object can be elevated to a statement with topological or spectral meaning.

The two presentations carry different symmetries. The additive presentation is preserved by translations of the summation variable; the multiplicative presentation is preserved by scalings of the multiplicative variable. More sharply, the two presentations are supported by free structures of different rank. On the additive side, the monoid $(\mathbb{N}, +)$ is generated by the single element 1; its rank is 1. On the multiplicative side, the monoid $(\mathbb{N}, \cdot)$ is freely generated by the set $\mathbb{P}$ of primes; its rank is $|\mathbb{P}| = \infty$. Consequently, although the two presentations agree numerically, there is no structural identity between the objects that carry them. This rank gap is not a technical artefact. It is the exact place where the numerical equality of the two presentations and their structural non-identity part ways, and it is the place at which we shall locate the content of RH.

It is important to say at the outset what is *not* claimed. The rank gap does not by itself imply RH, and does not by itself refute anything. It locates a tension precisely: the two presentations agree as numbers, differ as structures, and the canonical map relating them is not an equivalence. The question that organises this paper is what the failure of that equivalence measures, and what replaces it.

1.2 The structural question and why the naive reading fails

Given the above, the natural structural question is whether the two sides can be related by a map, and if so, what the failure of that map to be an equivalence measures. The most direct reading — "the two sides are homotopy equivalent as chain complexes, RH being the vanishing of the obstruction class" — is the reading that a first structural investigation would propose. A substantial part of this paper is devoted to showing that this reading is **too strong**, and to identifying what should replace it.

Specifically, we show in §4 that there is a canonical chain map between the two sides, given by the fundamental theorem of arithmetic, and that this map is injective but not surjective. Consequently no chain map between the two sides can induce an isomorphism on homology, because the free abelian groups $H_0(C_{\text{add}})$ and $H_0(C_{\text{mul}})$ have ranks 1 and ∞ . The requirement of homotopy equivalence is therefore impossible, and any programme that seeks an "object" at the level of complexes is seeking an object that cannot exist. This is a negative result, but a productive one: it tells us that the correct level is not the level of complexes but the level of **pairings** and of **spectra**.

Once the level is corrected, the programme decomposes into the pieces below (§3–§8): two premises that fix the line on which the content lives, theorems about the objects, and one sharply stated existence problem that §6.6–§6.7 localise. We state the corrections here so that the reader may follow the logic of the whole paper from the outset.

1. The two sides are not two complexes without a common base; they are **one pairing** expressed twice, and the identity that expresses this is the Riemann–von Mangoldt explicit formula [2] (§5).
2. The obstruction δ_ζ is **not** undefined; it admits a well-defined cumulative form $\delta_\zeta(x) = \psi(x) - x + \log 2\pi + 1/2 \log(1 - x^{-2})$, and the statement " $\delta_\zeta = 0$ " is not equivalent to RH but is an unconditional identity. The RH content lies in a **positivity** condition on the pairing, not in the vanishing of δ_ζ (§5, §6).

3. Given the spectrum, the operator carrying the RH content is **unique** up to unitary equivalence, its density is locked by the zero-counting function, and self-adjointness is load-bearing: removing it destroys uniqueness (§6).
4. What remains is **existence** — the Hilbert–Pólya statement — and this is equivalent to two computable positivity criteria, those of Li [11] and of Weil [12, 13, 14], rather than being a notion without content (§6, §9). It is not treated here as an "honesty gap"; it is treated as the one item the structural route leaves open, and its equivalent form is computed.
5. The critical line is **not** a convention, and the value $1/2$ is not inserted through the functional equation and read back. The same value is returned independently by the reflection symmetry of the completed zeta function (analytic route) and by the self-normalisation exponent of the quadratic Gauss sum, $|G(p)| = \sqrt{p}$ (arithmetic route, using no analytic input); the two coincide (§7.7).
6. The discrete reciprocity symbol ± 1 and the continuous local angle θ_p are **not** objects of different type; they are the same object across a normalisation seam, and the seam is pinned at $1/2$ — the seam lift (§7.8).
7. The one **open** item is *localised*, not left unlocated: the operator route is exhibited in two directly comparable structures where the analogue of RH is a theorem, and the existence statement ES is equivalent to a single condition — positive-definiteness at the archimedean place — that is free in those structures and conditional here (§6.6).
8. The same open item is *localised a second time*, in a dynamics: the de Bruijn–Newman backwards heat flow H_t (with $H_0 = 1/8\xi(1/2 + iz/2)$) has all real zeros exactly for $t \geq \Lambda$, and $\text{RH} \Leftrightarrow \Lambda = 0$. The zero dynamics is the gradient flow of a logarithmic Hamiltonian with monotone non-increasing energy $E(t) = \sum_{j \neq k} |x_j - x_k|^{-2}$, so that least dissipation forces the configuration to an equal-spacing equilibrium. The flow acts on the real variable and does not by itself select a vertical line: the *location* of the baseline is not an output of the dynamics but the value $1/2$ already fixed by the spatial route of §7.7, which the least-dissipation optimum therefore **meets** rather than derives. Unconditionally $0 \leq \Lambda \leq 0.22$ [15, 16], and the residual content is the single scalar $\Lambda = 0$ — the *time face* of the baseline (§6.7).

The methodological point of this organisation is that each correction converts a statement that *looked* like a gap into either a theorem or an explicitly computable condition. A structural programme that stops at "there is no object" has left work undone; the present one carries each piece to the level at which it becomes a definite statement.

1.3 Contributions

The paper's contributions are the following, each stated with the section in which it is established.

1. **An explicit chain map.** We construct $\varphi(n) = \sum_p v_p(n)[p]$ between the two sides and prove that it is a homomorphism (Proposition 4.1), an injection (Proposition 4.2), and not a surjection (Proposition 4.3) (§4.2).
2. **A rank-gap theorem.** We prove that no chain map between the two sides induces a homology isomorphism, hence that homotopy equivalence of the two complexes is impossible (Theorem 4.5). This converts "there is no object" from a diagnosis into a theorem (§4.3).
3. **A well-defined obstruction functional.** We replace the divergent pointwise series $\sum_p (\log p) \theta_p(s)/\pi$ by the cumulative functional $\delta_{-\zeta}(x) = \psi(x) - x + \log 2\pi + 1/2 \log(1 - x^{-2}) = -\sum_p x^{\wedge} p/p$, prove it is well-

defined for every $x > 1$, and verify that its prime side and zero side agree numerically (Theorem 5.1, Result 5.3) (§5).

4. **A corrected reading of $\delta_\zeta = 0$.** We show that the vanishing of the cumulative obstruction is an identity rather than a criterion, and that the RH content is carried by a positivity condition on the pairing (Proposition 5.2) (§5.4).
5. **A spectral uniqueness theorem.** We prove that self-adjoint operators with coinciding spectra (with multiplicity) are unitarily equivalent (Theorem 6.1); we exhibit a Jordan-block counterexample showing that self-adjointness is load-bearing (Proposition 6.3); and we prove that the eigenvalue counting function is locked to the Riemann–von Mangoldt density (Proposition 6.6) (§6).
6. **A reduction to a single open statement.** We prove the programme reduces to the single existence statement ES, every other link of the chain being settled as a theorem or proposition (Theorem 6.10) (§6.5).
7. **An analytic-continuation correspondence.** We record the completed-zeta functional equation $\xi(s) = \xi(1-s)$ [1, 8] and use it to show that the cumulative obstruction is defined on the whole range of interest, the critical line being the fixed line of the symmetry (§7).
8. **A structural baseline.** We prove the critical line is over-determined: the value $1/2$ is the unique reflection axis of ξ (analytic route, Proposition 7.11) and the exact self-normalisation exponent of the quadratic Gauss sum $|G(p)| = \sqrt{p}$ (arithmetic route, Proposition 7.12), the two derived independently of one another (Corollary 7.13) (§7.7).
9. **A seam lift.** We show that the discrete quadratic-reciprocity character lifts to the continuous local Euler phase θ_p , and that the lift is pinned at the baseline $1/2$ (Proposition 7.15), which closes the apparent type mismatch between the discrete symbol and the continuous angle (§7.8).
10. **A topological bridge.** We construct the combinatorial Chern character $ch = (\text{rank}, c_1)$ with $c_1 = \chi = V - E$, prove the discrete Gauss–Bonnet identity (Proposition 8.2) and the ring-homomorphism properties (Proposition 8.3), and identify it as the topological dual of the explicit-formula bridge (§8).
11. **A computable positivity form of the remaining statement.** We verify numerically that the Li coefficients λ_n are positive for $n \leq 20$ and compare λ_1 with its analytic value; we record Weil positivity as the equivalent analytic form [12, 13, 14] (§6.4, §9).
12. **A complete ledger of gaps.** We account for every item recorded as a "gap" in an earlier treatment, classifying each as closed, restated, or explicitly open (§10.11).
13. **Reproducibility.** All constructions are implemented in public single-file scripts with self-tests, and every numerical claim in the paper is produced by a documented protocol (§11, Appendix A–B).
14. **A localisation of the open item.** We exhibit the operator route in two second structures — function fields (Weil 1948) and compact hyperbolic surfaces (Selberg 1956) — in which the analogue of RH is a theorem because self-adjointness is free, a positive-definite pairing being supplied by a polarisation (Theorem 6.14, Theorem 6.15). We compute the discrete zero positions on the function-field structure (123 elliptic curves; every Frobenius eigenvalue with $|\alpha| = \sqrt{p}$ and every zero on $\text{Re } s = 1/2$, worst deviation 1.1×10^{-16}) and compare them with the arithmetic zeros. We then prove that G6 is equivalent to the positive-definiteness of the arithmetic pairing at the archimedean place (Theorem 6.17), localising the single open statement to one named ingredient (§6.6).

15. **A dynamical reduction of the open item to a bracketed scalar.** We show that G6 reduces further, via the de Bruijn–Newman heat flow, to the single assertion $\Lambda = 0$ (Theorem 6.22, Theorem 6.23, Proposition 6.26), with the unconditional two-sided bracket $0 \leq \Lambda \leq 0.22$ of Rodgers–Tao [15] and Polymath 15 [16]. We exhibit the zero dynamics as the gradient flow of a logarithmic Hamiltonian with monotone non-increasing energy (Theorem 6.24), read the energy as a reciprocal (harmonic-mean) functional whose ground state is fixed by the genus/density datum, and verify numerically that the flow relaxes monotonically to the equal-spacing ground state while the arithmetic state sits *above* it (reciprocal efficiency $\eta = 0.772$; unfolded-spacing variance 0.103 against an equilibrium value of 0) — exactly the input used in [15] to prove $\Lambda \geq 0$ (§6.7). The route is deliberately read in reverse: it forces the existence of the baseline, not its value. We formalise the two quantities this reading rests on: the dissipation work $W = \int E dt = 1/4 \Delta H$ (Definition 6.29, Proposition 6.30) and the reciprocal efficiency $\eta = E_{\text{lat}}/E_{\text{(act)}} \in (0,1]$ (Definition 6.31, Lemma 6.32, Lemma 6.33, Proposition 6.34), which pin the least-dissipation route to the genus potential (§6.8).

1.4 Central questions

The paper is organised around four questions, each corresponding to one of the four structural levels that the construction passes through.

- **Q1 (existence of comparison).** Can the two sides be made structurally comparable at all, given that their free structures have different rank?
- **Q2 (level of identity).** At which level — complexes, pairings, or spectra — are the two sides identical, and what object carries the RH content?
- **Q3 (uniqueness and rigidity).** Given the RH content as a spectral condition, is the corresponding operator unique, and is its density free?
- **Q4 (boundary).** Where exactly does the programme stop, and is what remains a genuine gap or an item with computable content?

The paper answers Q1 in §4 (by the chain map and the rank-gap theorem), Q2 in §5 (the pairing identity), Q3 in §6 and §7 (the spectral uniqueness theorem, the density lock, the structural baseline of §7.7, and the dynamical critical parameter of §6.7), and Q4 in §9 and §10 (the reduction and the positivity criteria).

1.5 Organisation of the paper

§2 fixes notation and recalls the relevant background: the two sides and their rank gap, chain complexes and chain maps, the explicit formula, the spectral theorem and the Hilbert–Pólya programme, the positivity criteria of Li and Weil, arithmetic topology, and the Chern character. §3 states the problem formally, with the definitions of the two sides and of the obstruction, and shows that the formulation is not vacuous. §4 constructs the chain map and proves the rank-gap theorem. §5 re-defines the obstruction functional in a well-defined cumulative form and establishes the pairing identity. §6 proves the spectral uniqueness theorem, the load-bearing counterexample, and the density lock, reduces the programme to a single existence statement with two positivity forms (§6.4–§6.5), localises that statement first to archimedean positive-definiteness via the operator route in two second structures (§6.6), then, dynamically, to the bracketed scalar $\Lambda = 0$ of de Bruijn–Newman, showing that the least-dissipation flow necessarily possesses a baseline which coincides with the value $1/2$ already fixed by §7.7 (the flow cannot select the line itself) (§6.7), and formalises the dissipation work and the reciprocal efficiency on which that reading rests, tying the least-dissipation route to the genus potential (§6.8). §7 gives the analytic-continuation correspondence, establishes the critical line as a structural output over-determined by an analytic and an arithmetic route (§7.7), and closes the discrete/continuous seam at that

baseline (§7.8). §8 constructs the Chern-character bridge. §9 verifies all constructions under one protocol and cross-checks them against an independently written probe. §10 discusses the three-level mapping, the correspondence with the Hilbert–Pólya programme, the relation to trace formulas and to the probability-cloud reading of the zeros, and the limitations and falsifiability of the construction. §11 gives the methods, with the complete core code and its creation logic. §12 concludes. References, appendices and declarations follow.

2 Mathematical Background and Notation

This section fixes notation and recalls the material on which the construction rests. Nothing here is new; the purpose is to place the reader in a position to read §3–§8 without consulting external sources, and to make explicit the precise form in which each classical ingredient enters the argument.

2.1 The two sides and their rank gap

Let $\mathbb{P}$ denote the set of primes and let $\mathbb{N} = \{1, 2, 3, \dots\}$. For $\operatorname{Re} s > 1$,

$$\zeta(s) = \sum_{n \geq 1} n^{-s} = \prod_{p \in \mathbb{P}} (1 - p^{-s})^{-1}. \quad (2.1)$$

We call the first expression the **additive side** and the second the **multiplicative side** of the zeta function. The equality in (2.1) is a consequence of the fundamental theorem of arithmetic [10], which states that every $n \geq 1$ admits a unique factorisation

$$n = \prod_{p \in \mathbb{P}} p^{v_p(n)}, \quad v_p(n) \in \mathbb{Z}_{\geq 0}, \quad v_p(n) = 0 \text{ for all but finitely many } p, \quad (2.2)$$

where v_p denotes the p -adic valuation.

Definition 2.1 (rank of a free structure). For a commutative monoid M that is free on a generating set S , the **rank** of M is the cardinality $|S|$. For a free abelian group $G \cong \mathbb{Z}^{(S)}$, the rank is $|S|$.

With this convention:

- the additive monoid $(\mathbb{N}, +)$ is generated by the single element 1, hence has rank 1;
- the multiplicative monoid $(\mathbb{N}, \cdot)$ is freely generated by $\mathbb{P}$, hence has rank $|\mathbb{P}| = \infty$.

Observation 2.2. The equality (2.1) identifies two presentations carried by free structures of different rank. The identification is accomplished by the map supplied by (2.2), which is injective but not surjective (§4). The numerical equality of the presentations therefore coexists with the structural non-identity of the objects carrying them.

It is worth stating explicitly what Observation 2.2 does and does not say.

- It does **not** say that the zeta function is "two things". There is one meromorphic function ζ ; what is multiple is the presentation.
- It does **not** say that the two sides are unrelated. They are related by an explicit map, which we construct in §4.
- It **does** say that any statement of the form "the two sides are the same object" must specify the categorical level at which sameness is asserted. At the level of complexes, sameness fails (§4.3). At the level of pairings, sameness holds (§5.3). These are not contradictory; they are statements about different levels.

The remainder of the paper is, in effect, an unpacking of the last bullet.

2.2 Chain complexes, chain maps, and homology

We work over $\mathbb{Z}$ unless stated otherwise. All complexes considered are bounded below and of finite type.

Definition 2.3 (chain complex and chain map). A chain complex $(C_\bullet, \partial_\bullet)$ is a sequence of abelian groups C_k together with homomorphisms $\partial_k : C_k \rightarrow C_{k-1}$ satisfying $\partial_{k-1} \circ \partial_k = 0$. A chain map $f : (C_\bullet, \partial) \rightarrow (D_\bullet, \partial')$ is a family of homomorphisms $f_k : C_k \rightarrow D_k$ satisfying $\partial'_k \circ f_k = f_{k-1} \circ \partial_k$ for every k . Its homology is the family $H_k(C) = \ker \partial_k / \operatorname{im} \partial_{k+1}$, and f induces $f_* : H_k(C) \rightarrow H_k(D)$.

Definition 2.4 (homotopy equivalence). Two complexes are homotopy equivalent if there exist chain maps $f : C \rightarrow D$ and $g : D \rightarrow C$ with $g \circ f \cong \operatorname{id}_C$ and $f \circ g \cong \operatorname{id}_D$ (chain homotopic to the respective identities). A homotopy equivalence induces isomorphisms $f_* : H_k(C) \rightarrow H_k(D)$ for all k .

The content of Definition 2.4 that matters here is the last sentence: **homotopy equivalence implies homology isomorphism**. The contrapositive is the engine of §4: to prove that no homotopy equivalence exists, it suffices to prove that no homology isomorphism exists. And for free abelian groups, homology isomorphism is controlled by rank, which is an isomorphism invariant:

Fact 2.5. Free abelian groups of different (possibly infinite) rank are not isomorphic.

We shall use Fact 2.5 with ranks 1 and ∞ . (For infinite rank the statement requires the standard fact that a free abelian group determines its rank up to the cardinality of a basis; this is the content of the invariant basis number property for free abelian groups.)

The two complexes used in this paper. For the additive side we use the 1-skeleton of the standard ordering:

$$C_-(\text{add}) = C_-(\text{add})(N): \quad \text{vertices } \{1, \dots, N\}, \quad \text{edges } \{(n, n+1) : 1 \leq n < N\}.$$

This is a path graph. Its homology is

$$H_0(C_-(\text{add})) = \mathbb{Z} \quad (\text{the graph is connected}), \quad H_1(C_-(\text{add})) = 0 \quad (\text{a tree has no cycle}). \quad (2.3)$$

These values are independent of N . The additive side is thus modelled by a complex whose degree-zero homology is free of rank 1.

For the multiplicative side we use the free abelian group on the primes,

$$C_{\text{mul}} = \mathbb{Z}^\wedge((\mathbb{P})), \quad H_0(C_{\text{mul}}) = \mathbb{Z}^\wedge((\mathbb{P})), \quad (2.4)$$

of rank $|\mathbb{P}| = \infty$. The contrast between (2.3) and (2.4) is the rank gap of §2.1, now expressed homologically.

Remark 2.6 (on the choice of model for the additive side). The path graph is one of many complexes that model the additive side; what matters for §4 is only that its degree-zero homology be free of rank 1. Any connected graph (indeed any complex with $H_0 = \mathbb{Z}$) would serve. The path graph is chosen because it is the 1-skeleton of the natural ordering of $\mathbb{N}$, and because its higher homology vanishes, so that no accidental cycles contaminate the comparison. The numerical probe of §9.3 verifies that replacing the path graph by other connected graphs on $\{1, \dots, N\}$ leaves $b_0 = 1, b_1 = 0$ for all tested N .

2.3 The explicit formula

The Riemann–von Mangoldt explicit formula [2] relates the distribution of primes to the zeros of ζ :

$$\psi(x) := \sum_{n \leq x} \Lambda(n) = x - \sum_{\rho} (x^\rho)/(\rho) - \log 2\pi - 1/2 \log(1 - x^{-2}), \quad x > 1, \quad (2.5)$$

where Λ is the von Mangoldt function ($\Lambda(n) = \log p$ if n is a power of the prime p , and 0 otherwise), and ρ runs over the non-trivial zeros of ζ , taken in conjugate pairs so that the sum is real. The formula is valid for $x > 1$ with the convention that if x is a prime power the left-hand side is reduced by $1/2\Lambda(x)$; alternatively one may restrict to x not a prime power, as we do throughout.

Why (2.5) is the right instrument here. The left-hand side of (2.5) is computed from the primes alone; the right-hand side is computed from the zeros alone. The formula therefore asserts the equality of two functions that are assembled from the two structures of §2.1, and it does so *without* requiring the two structures to be isomorphic. This is the precise sense in which (2.5) is a **pairing identity**: it evaluates both sides against a common test object, the real parameter x , rather than matching the structures to one another. A pairing needs only a common domain on which to be evaluated; it does not need an isomorphism of carriers. This distinction is the reason the rank gap, which is fatal at the level of complexes, is inoperative at the level of pairings (§5.3).

The explicit formula is the prototype of the general spectral–geometric duality expressed by trace formulas [3]. In the language of that duality, the term x is the main term, and $-\sum_{\rho} x^{\rho}/\rho$ is the oscillatory tail whose discrete source is the zero set.

2.4 Spectral theorems and the Hilbert–Pólya programme

The Hilbert–Pólya programme seeks a self-adjoint operator whose spectrum is the set of non-trivial zero ordinates: writing the non-trivial zeros as $\rho_n = 1/2 + i\gamma_n$, one asks for a self-adjoint H with

$$\text{spec}(H) = \{\gamma_n\}_{n \geq 1}.$$

If such an operator exists, then its spectrum is real by the spectral theorem, whence every γ_n is real and every non-trivial zero lies on the critical line — that is, RH follows. The programme's physical motivation, and the statistical evidence for it, are reviewed by Odlyzko [17], Berry and Keating [18] and Connes [4].

Two logically distinct halves must be distinguished, and this distinction is one of the organising ideas of the present paper.

- The **existence half** asks: is there a self-adjoint operator with spectrum $\{\gamma_n\}$? This is open [7].
- The **uniqueness half** asks: given a spectrum, how many self-adjoint operators realise it? This is a theorem, and it is the content of §6.2.

Conflating the two halves is a common source of confusion. Separating them is what allows the programme to be reduced to a single clearly stated problem (§6.5): the uniqueness half is settled by the spectral theorem, so no effort is wasted in asking whether the operator is "free"; the existence half is the one item that remains, and it is shown in §6.4 to be equivalent to a positivity condition with numerical content.

The spectral theorem in the form we need is the following.

Theorem 2.7 (spectral theorem, discrete self-adjoint form). Let H be a self-adjoint operator on a separable Hilbert space with purely discrete spectrum. Then H is unitarily equivalent to the multiplication operator by the independent variable on a direct sum $\bigoplus_j L^2(\mathbb{R}, \mu_j)$, and its unitary equivalence class is determined by the spectral measure class together with the multiplicity function $m(\lambda) = \dim \ker(H - \lambda I)$.

We use Theorem 2.7 in §6.2, where the consequence is that self-adjoint operators with the same spectrum *with multiplicity* are unitarily equivalent. The qualifier "with multiplicity" is essential and is the reason the Jordan block of Proposition 6.3 is a counterexample when self-adjointness is dropped: the Jordan block $J = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$ has the same eigenvalues as $0_{(2 \times 2)}$ but different multiplicities in the sense of $\dim \ker(H - \lambda I)^k$.

The de Bruijn–Newman flow [15, 16, 19, 20]. A third equivalent form of RH is dynamical. With Φ the super-kernel of Definition 6.21 and

$$H_t(z) = \int_0^\infty e^{-(t+u^2)} \Phi(u) \cos(zu) \, du, \quad \text{so that} \quad H_0(z) = 1/8 \xi(1/2 + iz/2),$$

the family $\{H_t\}$ solves the backwards heat equation $\partial_t H_t = -\partial_z^2 H_t$, de Bruijn and Newman showed that there is a constant Λ with H_t real-rooted **exactly** for $t \geq \Lambda$; RH is the assertion $\Lambda \leq 0$. The two bounds

$$\Lambda \geq 0 \quad [15], \quad \Lambda \leq 0.22 \quad [16]$$

are unconditional, so $0 \leq \Lambda \leq 0.22$ and $\text{RH} \Leftrightarrow \Lambda = 0$. Moreover the zero dynamics is the gradient flow of a logarithmic Hamiltonian whose energy is monotone non-increasing [15]. This material is used in §6.7, where the flow is shown to possess a baseline as its least-dissipation equilibrium — the baseline $1/2$ already fixed by §7.7 rather than re-derived, since a flow in the real variable cannot select a vertical line — and the open item is thereby reduced to the bracketed scalar Λ .

2.5 Positivity criteria: Li and Weil

Two positivity conditions are known to be equivalent to RH, and they give the existence half of §2.4 computable content.

Li's criterion [11]. Define, for $n \geq 1$,

$$\lambda_n = \sum_{\rho} [1 - (1 - 1/\rho)^n], \quad (2.6)$$

the sum being over the non-trivial zeros in the usual symmetric (conditionally convergent) sense. Then

$$\text{RH} \Leftrightarrow \lambda_n \geq 0 \quad \text{for all } n \geq 1. \quad (2.7)$$

The first coefficient has the closed form

$$\lambda_1 = 1 + (\gamma)/(2) - 1/2 \log 4\pi \approx 0.023096, \quad (2.8)$$

where γ is the Euler–Mascheroni constant; this supplies an analytic target against which numerical partial sums can be compared. The generalisation of the criterion to arbitrary multisets and its relation to Weil's criterion are due to Bombieri and Lagarias [21].

Weil positivity [12, 13]. Let f be a suitable test function (say, in the Schwartz class with Fourier transform supported in a suitable range) and let $f \star \tilde{f}$ denote its autocorrelation, where $\tilde{f}(x) = \overline{f(-x)}$. Then

$$\text{RH} \Leftrightarrow W(f \star \tilde{f}) \geq 0 \quad \text{for all such } f, \quad (2.9)$$

where W is the Weil functional. In the noncommutative-geometric formulation of Connes [4] and Connes–Consani [14] this positivity is the trace-formula statement of the same condition; the archimedean place for which positivity is hardest to establish is analysed in [14].

How the two criteria enter this paper. They enter at §6.4, where the existence statement ES is shown to have, in addition to its original form, the two equivalent forms (2.7) and (2.9). The numerical work of §9 is confined to (2.7), because (2.7) reduces, on a finite range, to an elementary computation over a finite list of zeros; (2.9) is quoted as an equivalent criterion rather than independently reproved here, a limitation recorded explicitly in §10.4.

2.6 Arithmetic topology and the Chern character

Two further pieces of background are used in §8 as a bridge, dual in spirit to the analytic bridge of §2.3.

Arithmetic topology [22] draws an analogy between primes and knots, in which the quadratic reciprocity law plays the role of the linking number. The analogy supplies the only non-arbitrary "connection" between primes: whereas one may draw an arbitrary graph on the primes using any rule whatsoever, the reciprocity symbol imposes a canonical, arithmetic structure. We use this in §8.4 to contrast arbitrary connection rules

with the canonical one. The reciprocity datum is not only a contrast, however: its quadratic Gauss sum supplies the arithmetic route to the baseline of §7.7, and its normalisation supplies the seam lift of §7.8, by which the continuous local Euler-factor phase of §5.1 is identified as the normalised image of the discrete character.

The Chern character [23, 24, 25] is the ring homomorphism

$$\text{ch} : K^0(X) \rightarrow H^{\wedge}(\text{even})(X, \mathbb{Q})$$

that carries the K-theory class of a complex vector bundle to a cohomology class; for a bundle E with Chern roots y_i (i.e. $c(E) = \prod_i (1 + y_i)$), $\text{ch}(E) = \sum_i e^{\wedge}(y_i)$. In §8 we use the combinatorial form that this takes on a finite graph, where it reduces to the pair

$$\text{ch} = (\text{ch}_0, \text{ch}_1) = (\text{rank}, c_1), \quad c_1 = \chi = V - E, \quad (2.10)$$

with χ the Euler characteristic. The motivation for introducing it here is dualism: the explicit formula is an *analytic* bridge between the two sides, whereas the Chern character is a *topological* one. Both bring the two sides into a common language — the former into a common value, the latter into a common class — and both are therefore instances of the "identification at a weaker level" that §5.1 promises. The index-theoretic context in which the Chern character and the Todd class reproduce the index of an elliptic operator is given by Atiyah and Singer [26] and, for the differentiable Riemann–Roch theorem, by Atiyah and Hirzebruch [24].

2.7 Positioning against related work

Table 2.1 records how this paper differs from the principal related works.

Work	Content	Relation to this paper
Riemann (1859) [1]; von Mangoldt (1895) [2]	Definition; explicit formula	We read the explicit formula as the <i>pairing identity</i> between two sides and re-derive the obstruction from it (§5)
Hilbert–Pólya (Odlyzko [17]; Berry–Keating [18]; Connes [4])	Operator whose spectrum is the zeros	We settle the <i>uniqueness</i> half given the spectrum; existence remains open (§6.2, §6.4)
Weil (1952, 1972) [12, 13]; Connes–Consani (2021) [14]	Positivity $\Leftrightarrow$ RH	We place positivity as the <i>existence half</i> and compute an equivalent form (§6.4)
Li (1997) [11]	$\lambda_n \geq 0 \Leftrightarrow$ RH	We compute λ_n as a <i>checkable</i> substitute for "no object" (§6.4, §9.2)
Bombieri–Lagarias (1999) [21]	Generalisation of Li's criterion; Li follows from Weil	We use Li's criterion as the numerically accessible form (§6.4, §9.2)
Montgomery (1973) [5]	Pair correlation of zeros (GUE statistics)	Statistics of the zeros, not their structural provenance (§10.6)
Selberg (1956) [3]	Trace formula	We read the explicit formula as the number-theoretic instance of the same paradigm (§10.3)
Morishita (2012) [22]	Primes as knots; reciprocity as linking	We use reciprocity as the canonical connection (§8.4) and, through its quadratic Gauss sum, as the arithmetic route to the baseline (§7.7–§7.8)
Chern (1946) [23]; Atiyah–Hirzebruch (1959) [24]; Atiyah–Singer (1963) [26]	Chern character; index theorem	We use the graph-level reduction of ch as the topological bridge (§8)
Platt–Trudgian (2021) [6]	Verification to height 3×10^{12}	Numerical verification; we do not

		compete with it, we restructure the statement
de Bruijn (1950) [19]; Newman (1976) [20]; Rodgers–Tao (2018) [15]; Polymath 15 (2019) [16]	The de Bruijn–Newman constant Λ ; RH $\Leftrightarrow \Lambda \leq 0$; $0 \leq \Lambda \leq 0.22$	We use the heat flow to localise the open item to the scalar $\Lambda = 0$; the least-dissipation equilibrium it forces is the baseline 1/2 already fixed in §7.7, not a re-derivation of it (§6.7)
This paper	Chain map + well-defined obstruction + spectral uniqueness + two bridges	Turns "two incompatible complexes" into "one pairing + one open existence statement with computable positivity"

2.8 A worked numerical example of the pairing

To fix the abstract discussion of §2.3 and to prepare §5, we compute the two sides of (2.5) at a single point. Take $x = 100$. On the prime side, the exact evaluation gives $\psi(100) = 94.0453$, and the main-term correction yields

$$\delta_{\zeta}(100) = \psi(100) - 100 + \log 2\pi + 1/2 \log(1 - 100^{-2}) = -4.116862.$$

On the zero side, summing the first 20 conjugate pairs,

$$-\sum_{\gamma > 0} \left((100^{1/2 + i\gamma})/(1/2 + i\gamma) + (100^{1/2 - i\gamma})/(1/2 - i\gamma) \right) = -2.480154.$$

The two values differ by a relative residual of 1.64×10^{-2} . The residual is not an error in the identity; it is the truncation of the zero sum, and it is instructive to see why. The function δ_{ζ} is the difference between a step-like counting function and a smooth main term; the 20 zeros retained capture the oscillations of wavelength comparable to $\log(100)$ and longer, while the omitted zeros contribute oscillations of higher frequency. As x grows, the same retained zeros approximate the oscillating tail proportionally better, and the relative residual falls to 1.06×10^{-4} at $x = 10^5$ (§5.5). This example also shows why the *cumulative* form, and not the pointwise series, is the correct representation: the cumulative difference is finite at every $x > 1$, whereas the pointwise analogue diverges on the critical line (§5.1).

2.9 Functorial reading of the two sides

It is worth isolating the general principle behind the constructions of §3–§4, because doing so explains which of the paper's statements are about the zeta function specifically and which are about a wider class of objects.

The additive presentation is a sum over the free commutative monoid on one generator; the multiplicative presentation is a product over the free commutative monoid on the set of primes. Both presentations are obtained from a monoid by a standard construction: from $(\mathbb{N}, +)$ one forms the generating series $\sum_n n^{\cdot}(-s)$, and from $(\mathbb{N}, \cdot)$ one forms the Euler product $\prod_p (1 - p^{\cdot}(-s))^{-1}$. The map between the two monoids is the factorisation map ϕ , and the equality of the two presentations is the statement that ϕ is compatible with the two generating constructions.

This is a functorial statement, and it is exactly at the functorial level that the rank gap appears. A functor that carries a free monoid to a generating series does not, in general, preserve rank: it carries a rank-one monoid and an infinite-rank monoid to two series that happen to agree. The agreement is a theorem about the objects (the fundamental theorem of arithmetic [10]); the failure of rank preservation is a statement about the functor. Three consequences follow, and they organise the whole paper.

1. **The agreement is not a structural identity.** Two presentations of one function need not be two presentations of one structure; the rank gap of §2.1 is the measure of the difference.
2. **The right comparison is at the level of evaluation, not of carriers.** A functor-landscape comparison that is blind to rank must be replaced by a comparison of values, i.e. a pairing (§5).
3. **The residual content is spectral.** What the functor does *not* determine — the placement of the zeros — is exactly what remains to be provided, and it is provided, if at all, by an operator (§6).

This functorial reading also makes precise why nothing in §4 depends on analysis. The rank gap is a statement about two free monoids and the map between them; it would hold verbatim for any pair of generating constructions related by an injective non-surjective map between free monoids of different rank. The zeta function enters only when one asks what the residual content *is* — and that question is answered by the explicit formula and the spectral statement, not by the rank comparison.

Finally, the functorial reading suggests a natural question, recorded here and returned to in §10.9: is the pair of constructions (generating series, Euler product) the *only* pair of functors under which the two sides agree, or are there others? The two bridges of §5 and §8 — the analytic bridge to a common value and the topological bridge to a common class — are two known common targets, and the question of whether the list is exhaustive is a structural question that does not depend on RH being true or false.

3 Problem Formulation

With the two sides and their rank gap fixed (§2.1), we can now state the problem precisely. The purpose of this section is to define the two objects whose interplay drives the whole paper — the two-sided structure and its obstruction functional — and to state, in a form that can be tested, what it would mean for the two sides to "agree". A formulation is worth writing down carefully here because the earlier, more naive formulation ("find a chain map whose vanishing obstruction class proves RH") is, as §4 shows, impossible; the corrected formulation is the one developed in §4–§8.

3.1 Definition of the two sides

Definition 3.1 (two sides). Let $C_{\text{add}} = C_{\text{add}}(\mathbb{N})$ be the path-graph chain complex on the vertex set $\{1, \dots, N\}$ with edges $(n, n + 1)$, and let $C_{\text{mul}} = \mathbb{Z}^{\wedge}(\mathbb{P})$ be the free abelian group on the primes, regarded as a complex concentrated in degree 0. The additive side carries the Dirichlet series $\sum_{n \geq 1} n^{-s}$; the multiplicative side carries the Euler product $\prod_p (1 - p^{-s})^{-1}$.

Two remarks on Definition 3.1 are necessary for what follows.

First, the complex C_{mul} is taken to be concentrated in degree 0. This is a modelling choice, not a theorem: it says that the multiplicative structure of $\mathbb{N}$ is carried by the *generators* (the primes), with the exponents as coefficients. If one instead modelled the multiplicative side by a complex with nontrivial higher groups, the rank of H_0 would be unchanged, and the rank-gap theorem of §4.3 would be unaffected, because that theorem uses only H_0 .

Second, we do not claim that C_{add} and C_{mul} are "the" complexes attached to the two sides. They are minimal models chosen so that the comparison is honest: the additive model records the ordering of $\mathbb{N}$, and the multiplicative model records the factorisation. The construction of §4 then tests whether these minimal models can be related.

Definition 3.2 (factorisation map). For $n \geq 1$ write $n = \prod_p p^{v_p(n)}$ as in (2.2), and set

$$\varphi(n) := \sum_{p \in \mathbb{P}} v_p(n) [p] \in \mathbb{Z}^\wedge(\mathbb{P}), \quad (3.1)$$

where $[p]$ denotes the basis element of $\mathbb{Z}^\wedge(\mathbb{P})$ corresponding to the prime p . Extend φ linearly to the degree-zero chains of $C_-(\text{add})$, and set $\varphi = 0$ in positive degrees.

Thus φ is the map "send an integer to the vector of its prime exponents". It is the only map from $\mathbb{N}$ to $\mathbb{Z}^\wedge(\mathbb{P})$ that is natural with respect to the two monoid structures, in the sense that it is forced by the requirement that $p \mapsto [p]$ for prime p and that the map be additive. The object φ , together with its properties, is studied in §4.

3.2 Definition of the obstruction

Definition 3.3 (obstruction functional, cumulative form). For $x > 1$ define

$$\delta_- \zeta(x) := \psi(x) - x + \log 2\pi + 1/2 \log(1 - x^{-2}), \quad (3.2)$$

where $\psi(x) = \sum_{n \leq x} \Lambda(n)$. By the explicit formula (§2.3) this equals

$$\delta_- \zeta(x) = - \sum_{\rho} x^{\wedge \rho} / (\rho). \quad (3.3)$$

The left-hand form (3.2) is computed from the primes; the right-hand form (3.3) is computed from the zeros. The two agree exactly (Theorem 5.1).

Remark 3.4 (why the cumulative form). It is important that Definition 3.3 is given in cumulative form. A natural earlier definition was the pointwise series

$$\Delta_- \zeta(s) := \sum_p ((\log p) \theta_p(s)) / (\pi), \quad \theta_p(s) = -\arg(1 - p^{-(1/2 - it)}), \quad s = 1/2 + it, \quad (3.4)$$

with θ_p the local Euler-factor phase at the critical line. As we show in §5.1, the series (3.4) diverges on the critical line, and the divergence is an artefact of the representation, not a property of the obstruction. Using the cumulative form (3.2) removes the artefact and makes $\delta_- \zeta$ a well-defined object on the whole range of interest. The distinction between (3.2) and (3.4) is thus not a technical convenience but a substantive correction: the obstruction is a global difference of counts, not a sum of local phases.

Lemma 3.5 (well-posedness of the cumulative obstruction). For every $x > 1$ that is not a prime power, the right-hand side of (3.2) is finite, and $\delta_- \zeta(x)$ is real.

Proof. $\psi(x)$ is a finite sum over the prime powers not exceeding x , hence finite. The terms x , $\log 2\pi$, and $1/2 \log(1 - x^{-2})$ are finite for $x > 1$. For the reality of (3.3): the non-trivial zeros occur in conjugate pairs $\rho, \bar{\rho}$ (equivalently, in pairs $1/2 \pm i\gamma$ under RH, and in conjugate pairs unconditionally), so each pair contributes $x^{\wedge \rho} / \rho + x^{\wedge \bar{\rho}} / \bar{\rho} = 2\text{Re}(x^{\wedge \rho} / \rho) \in \mathbb{R}$. $\square$

Lemma 3.5 is trivial, but it isolates exactly the property that (3.4) lacks: well-posedness on the critical line. The whole of §5.1 is a proof that (3.4) is not well-posed there, and §5.6 a comparison of three candidate representations.

3.3 What "the two sides agree" can mean

Before stating the problem in its three parts, we make explicit the several senses in which one could ask the two sides to "agree", because the choices are not equivalent and the paper's main correction consists in selecting the right one.

(A) Numerical agreement. The two presentations (2.1) define the same function. This holds, and is due to the fundamental theorem of arithmetic [10]. It is not a structural statement.

(B) Agreement of complexes (homotopy equivalence). One could ask whether C_{add} and C_{mul} are homotopy equivalent, with a putative "obstruction class" whose vanishing would be RH. This is the naive reading. §4.3 shows it to be impossible.

(C) Agreement of pairings. One could ask whether the two sides define the same pairing when evaluated against a common test object. This holds, and is the explicit formula (§5.3).

(D) Agreement of spectra. One could ask whether the two sides determine the same spectral data, i.e. whether there is a self-adjoint operator whose spectrum is the zero set. This is the Hilbert–Pólya statement; its uniqueness half is a theorem (§6.2), its existence half is open (§6.4). The open half is not unlocated: it is equivalent both to a positivity condition at the archimedean place (§6.6) and, dynamically, to the single scalar $\Lambda = 0$ of de Bruijn–Newman, bracketed unconditionally by $0 \leq \Lambda \leq 0.22$ (§6.7).

The paper's structural thesis is that (B) fails for a theorem-stated reason, that (C) is the correct replacement for (B), and that (D) is where the residual content lives. Formally:

$$(A) \Leftrightarrow (C) \quad (\text{both unconditional}), \quad (B) \text{ false}, \quad (D) \text{ open}.$$

3.4 Statement of the problem

The problem can now be stated in three parts, each addressed in a subsequent section.

- **Part A (the map).** Does there exist a chain map between the two sides, and is it an equivalence? §4 constructs the map (Propositions 4.1–4.3) and shows that it is not an equivalence, for a structural reason (Theorem 4.5).
- **Part B (the identity).** Since the two sides are not equivalent as complexes, at which level — if any — are they identical? §5 shows that they are identical as pairings, via the explicit formula (Theorem 5.1), and identifies the precise sense in which the vanishing of the obstruction is an identity rather than a criterion (Proposition 5.2).
- **Part C (the content).** What carries the RH content, and how much freedom does it have? §6 shows that the content is carried by a self-adjoint operator which, given its spectrum, is unique and has locked density (Theorem 6.1, Proposition 6.6); the only open item is existence, and it is equivalent to a positivity condition (Theorem 6.10). §6.6 localises that condition to the archimedean place and §6.7 reduces it further, dynamically, to the bracketed scalar $\Lambda = 0$.

3.5 Non-vacuity

A formulation of this kind is worth nothing unless it is non-vacuous: unless each part (A), (B), (C) can be decided, and unless the decision is not an artefact of the definitions. We record three elementary non-vacuity checks, each of which is implemented and reported in §9.

1. **The map of Part A exists and is non-trivial.** φ is not the zero map ($\varphi(2) = [2] \neq 0$) and not an isomorphism (Proposition 4.3). Hence the question "is it an equivalence?" has a definite negative answer rather than a vacuous positive one.
2. **The pairing of Part B is not an identity by construction.** The cumulative obstruction δ_ζ is defined *without* reference to the zeros (its left-hand form uses only ψ , x and elementary constants), and its agreement with the zero side (3.3) is a theorem, not a definition. Numerically, the two sides agree to a relative residual that decays with x (Result 5.3), which would be impossible if the agreement were built in.

3. **The spectral statement of Part C has content.** The uniqueness theorem is not vacuous because self-adjointness is load-bearing (Proposition 6.3 shows a same-spectrum non-self-adjoint operator that is not unitarily equivalent), and the density lock is not vacuous because the eigenvalue counting function is compared with an independently computed zero count (Result 6.3).

These three checks are the reason the corrected formulation (A)–(D) is not a reformulation for its own sake: each level is decided, and the decisions are independent of one another.

3.6 Why the corrected formulation is the minimal one

It is worth stating why the four senses (A)–(D) of "the two sides agree" are the natural ones, and why the corrected formulation does not introduce a level that the naive formulation omitted but rather *removes* a level that cannot be occupied.

The naive formulation asks for a level at which the two sides are isomorphic as carriers — sense (B). The difficulty with (B) is not that it is hard to decide but that it is *not occupied*: by Theorem 4.5 there is no chain map inducing a homology isomorphism, because the ranks differ. A level that does not exist cannot be searched, and recording its absence as an open problem misdescribes the situation. Sense (B) is therefore not a stronger version of (A) or (C); it is a level that the two sides never share.

Removing (B) leaves exactly (A), (C) and (D), and these three are ordered by strength in a way that is itself informative:

- (A) is the weakest and is unconditional: two presentations, one function [10].
- (C) is (A) restated at the level of evaluation: two sides, one value per test object (Theorem 5.5). It is unconditional and is what replaces (B).
- (D) is the strongest: it asserts the existence of an object whose spectrum is the zero set, which implies RH.

The corrected formulation is thus minimal in the following precise sense: it asserts at each of the three available levels exactly what that level can support, and it asserts nothing at the level that cannot be occupied. In particular it does not weaken the goal — (D) still implies RH — and it does not strengthen the intermediate steps beyond what is true (it does not claim (B)).

There is a second sense in which the formulation is minimal. Each of (A), (C), (D) is *decided* in this paper or is *explicitly open with a computable equivalent form*: (A) holds, (C) holds, and (D) is the single open item, whose equivalent forms are computed — a positivity condition at the archimedean place (§6.6) and, dynamically, a scalar bracketed by $0 \leq \Lambda \leq 0.22$ (§6.7). A formulation in which every level is either settled or computably testable leaves no room for a level whose status is merely unclear; and it is precisely such unclear levels that earlier treatments recorded as gaps. The minimality here is therefore not an aesthetic preference but the reason the paper has a definite answer to Q4: the boundary of the construction is a single, named, computably testable statement.

Finally, the formulation separates the *identity* from the *content* — the distinguishing feature of the corrected reading. Sense (A)/(C) is an identity: it holds regardless of where the zeros lie. Sense (D) is content: it is equivalent to their location. Confusing the two is what produced the earlier mis-description of $\delta_\zeta = 0$ as a criterion (§5.4). Keeping them apart is what allows §6 to isolate existence as the only item that carries information about RH.

4 The Chain Map and the Rank-Gap Theorem

This section constructs the map between the two sides and proves that it cannot be an equivalence. The logic is as follows. First (§4.1) we write down the map, which is forced by unique factorisation. Then (§4.2) we prove its three properties: it is a homomorphism, it is injective, and it is not surjective. Finally (§4.3) we deduce the rank-gap theorem, which is the reason the naive "homotopy equivalence" reading of the two sides must be abandoned, and which therefore motivates the change of level in §5. §4.4 records the numerical verification of all three properties.

4.1 Construction of the map

The map is the factorisation map of Definition 3.2,

$$\varphi(n) = \sum_p v_p(n)[p] : \mathbb{N} \rightarrow \mathbb{Z}^\wedge(\mathbb{P}).$$

Its construction uses nothing beyond the fundamental theorem of arithmetic [10]. Concretely:

$$\begin{aligned} \varphi(1) &= 0 \quad (\text{the empty factorisation}), & \varphi(2) &= [2], & \varphi(6) &= [2] + [12], \\ \varphi(12) &= 2[2] + [12], & \varphi(36) &= 2[2] + 2[12], & \varphi(210) &= [2] + [12] + [11] + [18]. \end{aligned}$$

In the numerical implementation (§11.3) we fix N and work in $\mathbb{Z}^\wedge(\mathbb{P}_N)$ with basis the primes up to N ; for $N = 200$ this basis has 46 elements, and the map is represented by a 46-dimensional integer vector for each $n \leq 200$.

Two features of the construction deserve comment.

First, φ is *canonical*: it is determined by the two monoid structures and by no further choice. Indeed, if one demands that a monoid homomorphism $\mathbb{N} \rightarrow \mathbb{Z}^\wedge(\mathbb{P})$ send each prime p to $[p]$, then it must send $n = \prod p^\wedge(v_p(n))$ to $\sum v_p(n)[p]$; there is nothing else it could do. This is what makes the three properties of §4.2 statements about the two sides rather than about an arbitrarily chosen map.

Second, φ is *not* an equivalence, and this is precisely what we wish to measure. The point of constructing the map is not the existence of a map — which is immediate from unique factorisation — but the failure of that map to be an isomorphism, which is a theorem (§4.3).

4.2 Three properties of the map

Proposition 4.1 (homomorphism). φ is a monoid homomorphism: $\varphi(ab) = \varphi(a) + \varphi(b)$ for all $a, b \geq 1$, where the left-hand product is in $(\mathbb{N}, \cdot)$ and the right-hand sum in $(\mathbb{Z}^\wedge(\mathbb{P}), +)$.

Proof. For each prime p , the p -adic valuation satisfies $v_p(ab) = v_p(a) + v_p(b)$; this is the additivity of the exponent in the product of two factorisations. Applying the definition (3.1) termwise,

$$\varphi(ab) = \sum_p v_p(ab)[p] = \sum_p (v_p(a) + v_p(b))[p] = \sum_p v_p(a)[p] + \sum_p v_p(b)[p] = \varphi(a) + \varphi(b). \quad \square$$

Proposition 4.2 (injectivity). φ is injective.

Proof. Suppose $\varphi(a) = \varphi(b)$ for some $a, b \geq 1$. Comparing coordinates in the basis $\{[p]\}_{(p \in \mathbb{P})}$, we obtain $v_p(a) = v_p(b)$ for every prime p . By the uniqueness part of the fundamental theorem of arithmetic [10],

$$a = \prod_p p^\wedge(v_p(a)) = \prod_p p^\wedge(v_p(b)) = b. \quad \square$$

Proposition 4.3 (non-surjectivity). φ is not surjective.

Proof. Consider the element $-[2] \in \mathbb{Z}^\wedge(\mathbb{P})$. If $-[2] = \varphi(n)$ for some $n \geq 1$, then comparing the coefficient of $[2]$ gives $v_2(n) = -1$, contradicting $v_p(n) \in \mathbb{Z}_+(\geq 0)$ for all p (Definition (2.2)). Since $-[2] \neq \varphi(n)$ for every n , we have $\varphi(\mathbb{N}) \subsetneq \mathbb{Z}^\wedge(\mathbb{P})$. $\square$

Remark 4.4 (the image of ϕ). The image $\phi(\mathbb{N})$ is exactly the submonoid of $\mathbb{Z}^{\wedge}(\mathbb{P})$ consisting of vectors with non-negative coordinates and finite support. Equivalently, $\phi(\mathbb{N}) = \bigoplus_p \mathbb{Z}_{\geq 0}$, the non-negative orthant of the free abelian group. The failure of surjectivity is therefore not a defect of the coordinate system but a statement about the two sides: the multiplicative side, regarded as a free structure, contains the group generated by the primes, whereas the image of $\mathbb{N}$ under factorisation contains only the monoid they generate. The difference is the presence of inverses, and it is exactly this difference that the rank gap of §4.3 measures in the infinite-rank case.

Propositions 4.1–4.3 are elementary. Their significance is that together they force the negative conclusion of §4.3, and that they are the *only* input that conclusion needs. No analytic input, no estimate, and no conjecture is used: the rank-gap theorem is a statement of algebra and unique factorisation alone.

4.3 The rank-gap theorem

We now assemble the elementary properties into the theorem that organises the rest of the paper.

Theorem 4.5 (rank gap). There is no chain map between C_{add} and C_{mul} inducing an isomorphism on homology. In particular, the two sides are not homotopy equivalent.

Proof. Suppose, for contradiction, that there is a chain map $f : C_{\text{add}} \rightarrow C_{\text{mul}}$ inducing an isomorphism $f_* : H_0(C_{\text{add}}) \rightarrow H_0(C_{\text{mul}})$. In degree 0 we have, by (2.3) and (2.4),

$$H_0(C_{\text{add}}) = \mathbb{Z}, \quad H_0(C_{\text{mul}}) = \mathbb{Z}^{\wedge}(\mathbb{P}) \cong \bigoplus_p \mathbb{Z}.$$

The left-hand group is free abelian of rank 1, the right-hand group free abelian of rank $|\mathbb{P}| = \infty$. By Fact 2.5, free abelian groups of different rank are not isomorphic; hence no isomorphism $H_0(C_{\text{add}}) \rightarrow H_0(C_{\text{mul}})$ exists, contradicting the assumption. Therefore no chain map induces a homology isomorphism.

For the second assertion: by Definition 2.4 and Fact 2.5, a homotopy equivalence would induce a homology isomorphism; since none exists, the two complexes are not homotopy equivalent. $\square$

Theorem 4.6 (rank gap, map-theoretic form). The canonical map ϕ of Definition 3.2 induces the zero map on positive-degree homology and an injective, non-surjective map $\phi_* : H_0(C_{\text{add}}) = \mathbb{Z} \hookrightarrow \mathbb{Z}^{\wedge}(\mathbb{P}) = H_0(C_{\text{mul}})$. Hence ϕ_* is not an isomorphism, consistently with Theorem 4.5.

Proof. Injectivity of ϕ_* on H_0 is Proposition 4.2 transported to homology (the degree-zero homology of a connected graph is generated by the class of a vertex, and ϕ is injective on vertices). Non-surjectivity is Proposition 4.3. The failure of isomorphism is then immediate. $\square$

Corollary 4.7 (the naive reading is impossible, not merely unproven). The statement "the two sides are homotopy equivalent, RH being the vanishing of the obstruction class" is false as stated. It is not a hard open problem; it is an impossible one.

Proof. By Theorem 4.5, no such equivalence exists. $\square$

Corollary 4.7 deserves a moment's emphasis, because it is the pivot of the paper. A structural programme that begins by asking for a homotopy equivalence between the two sides begins by asking for an object that cannot exist. This is a negative result, but a *productive* one: it tells us exactly where the programme must change level. The rank gap lives at the level of complexes; it disappears at the level of pairings, because a pairing does not require the two structures to be isomorphic, only to be evaluated against a common test object (Theorem 5.1). And it reappears, in a different and more tractable form, at the level of spectra, where it becomes the statement that the operator, if it exists, is unique with locked density (§6).

Corollary 4.8 (why the rank gap is the right locus). The rank gap is the precise obstruction to reading RH as a statement about the two complexes, and it is simultaneously the reason the pairing formulation of §5 is the correct one: the rank gap is a *structural* mismatch at the level of carriers, whereas the pairing identity of §5 is a statement at the level of evaluation, where carriers need not match.

Remark 4.5 (relation to the earlier treatment). In earlier treatments of this line of argument, the situation was described by saying that "there is no object" at the level of complexes. Theorem 4.5 upgrades that diagnosis to a theorem: the non-existence is not an imprecision of the search but a consequence of the rank of two free abelian groups. This is the first of the corrections listed in §1.2.

4.4 Numerical verification

Result 4.1 (chain map). Computed for $N = 200$ (basis of 46 primes):

- *homomorphism*: for all sampled pairs (a, b) with $ab \leq 200$, the identity $\varphi(ab) = \varphi(a) + \varphi(b)$ holds; e.g. $\varphi(6 \cdot 35) = \varphi(210) = [2] + [12] + [11] + [18] = \varphi(6) + \varphi(35)$.
- *injectivity*: no two of the 200 computed vectors $\varphi(1), \dots, \varphi(200)$ coincide.
- *non-surjectivity*: the vector $-[2]$ is not in the image; the check confirms that no $n \leq 200$ maps to it, and Proposition 4.3 shows this for all n .

Result 4.2 (rank gap). The rank of the additive free structure is 1; the rank of the multiplicative side is 46 (the number of primes up to $N = 200$), and in the limit $N \rightarrow \infty$ it is $|\mathcal{P}| = \infty$. Hence $\varphi_* : \text{Ho}(C_{\text{add}}) = \mathbb{Z} \rightarrow \text{Ho}(C_{\text{mul}}) = \mathbb{Z}^{\wedge}(\mathcal{P})$ is injective but is not an isomorphism, by Theorem 4.6.

Result 4.3 (structural invariance of the additive model). Replacing the path graph on $\{1, \dots, N\}$ by other connected graphs on the same vertex set leaves $b_0 = 1$ and $b_1 = 0$ for all tested N (the probe of §9.3 reports $b = (1, 0, 0)$ for $N = 10, \dots, 10^4$). The rank gap is therefore a property of the two sides and not of the particular additive model.



Fig. 1 | The two sides and the chain map. The homomorphism φ given by the fundamental theorem of arithmetic maps the multiplicative side into the free group on the primes; it is injective but not surjective, the two free structures having rank 1 and ∞ . The failure of surjectivity at the element $-[2]$ is displayed.

Result 4.4 (self-test). The three properties and the rank computation are exercised by the first three checks of the unified self-test of §9.2, which passes all checks (11/11 overall).

4.5 Why the rank gap is an invariant and not an artefact

The rank-gap theorem rests entirely on the fact that free abelian groups of different rank are non-isomorphic (Fact 2.5). It is worth pausing on this fact, because if it failed, the theorem would fail with it, and the rest of the paper would lose its pivot. The fact is not a convention about bases; it is a property of the groups themselves.

Lemma 4.9 (rank is well defined). If $\mathbb{Z}^{\wedge}(S) \cong \mathbb{Z}^{\wedge}(T)$ for sets S, T , then $|S| = |T|$.

Proof. For a free abelian group G , the cardinality of a basis is recovered from G without reference to any chosen basis: for finite rank, $\mathbb{Z}^n$ has 2^n homomorphisms to $\mathbb{Z}/2\mathbb{Z}$, so n is determined by the group; for infinite rank, the cardinality of a basis equals the cardinality of the set of homomorphisms $G \rightarrow \mathbb{Z}$ that are not divisible by any integer > 1 (equivalently, the rank equals $\dim_{\mathbb{Q}}(G \otimes \mathbb{Q})$ for finite rank, and $\dim_{\mathbb{Q}}(G \otimes \mathbb{Q})$ distinguishes finite from infinite; the cardinality statement for infinite rank follows by the same argument applied to $G \otimes \mathbb{Q}$ as a $\mathbb{Q}$ -vector space). In either case the rank is an isomorphism invariant. $\square$

Lemma 4.9 is the reason Theorem 4.5 is stated as a theorem rather than as an observation about the specific models C_{add} and C_{mul} : any model whose degree-zero homology is free of the appropriate rank falls under the same conclusion. Two further comments make the invariance concrete.

First, the conclusion is *independent of N* . For every finite N the multiplicative side has rank $\pi(N)$ (the number of primes up to N), and $\pi(N) \rightarrow \infty$ with N ; the additive side has rank 1 for every N . Thus the rank gap persists at every truncation and in the limit. There is no value of the cut-off at which the two sides become comparable as carriers; the mismatch is not a small-parameter effect that could be removed by taking N large.

Second, the conclusion is *independent of the model of the additive side*, provided only that its degree-zero homology be free of rank 1. This is verified empirically in Result 4.3, where replacing the path graph by other connected graphs leaves $b_0 = 1$ and $b_1 = 0$. The choice of the path graph is therefore a convenience, not a hidden hypothesis: the rank gap is a property of the two *sides*, expressed through whatever complexes model them.

The methodological point is the one recorded in §1.2. A negative structural result is only as good as the invariance of the quantity that produces it. Because the quantity here — rank of a free abelian group — is an isomorphism invariant of the homology, the negative conclusion is stable under every change of model that preserves homology, and it cannot be evaded by re-choosing the presentation of the additive side.

The consequence for the rest of the paper is the one anticipated in §1.2 and proved in Theorem 10.1: since the deciding content cannot live at the level of any homotopy invariant of the two sides, it must live at the *spectral* level. Section 6 exhibits it there; §6.6 localises it to the archimedean positive-definite pairing; and §6.7 reduces it, through the de Bruijn–Newman flow, to a single real constant inside a two-sided bracket. The rank gap of this section is therefore not only the obstruction to a topological reading but also the reason the argument is routed to spectra at all.

5 The Obstruction Functional and the Pairing Identity

Theorem 4.5 forces a change of level. If the two sides are not equivalent as complexes, the natural next question is whether they are identical at a weaker level, and if so which. The answer given in this section is that they are identical as **pairings**: the prime side and the zero side are two expressions of one quantity, and the identity expressing this is the explicit formula [2]. A crucial by-product (§5.1) is a correction of a natural but wrong definition of the obstruction functional: written as a pointwise series it diverges on the critical line, whereas written in cumulative form (§5.2) it is well-defined everywhere of interest. This is the second of the

corrections listed in §1.2: " δ_ζ is undefined" becomes " δ_ζ is well-defined, and it is the *vanishing* that is an identity rather than a criterion".

5.1 Why the pointwise series is the wrong representation

The local angle at a prime p on the critical line is the phase of the local Euler factor,

$$\theta_p(t) = -\arg(1 - p^{-(1/2 - it)}), \quad s = 1/2 + it. \quad (5.1)$$

Since $|p^{-(1/2 - it)}| = p^{-(1/2)} < 1$, the argument is a continuous, indeed real-analytic, function of t , and it is bounded:

$$\sup_{t \in \mathbb{R}} |\theta_p(t)| = \arcsin(p^{-(1/2)}), \quad \text{which equals } \pi/4 \text{ at } p = 2. \quad (5.2)$$

The bound (5.2) follows from the elementary fact that $|1 - w|$ with $|w| = p^{-(1/2)}$ is minimised, hence the argument maximised, when $1 - w$ is tangent to the circle of radius $p^{-(1/2)}$ about 1. Writing $w = p^{-(1/2)}e^{i\phi}$, the argument of $1 - w$ has absolute value $\arcsin(p^{-(1/2)}\sin\phi) \leq \arcsin(p^{-(1/2)})$.

The natural pointwise obstruction is then the series

$$\Delta_\zeta(s) = \sum_p ((\log p)\theta_p(s))/(\pi). \quad (5.3)$$

Its convergence behaviour depends on the half-plane.

Proposition 5.1 (convergence for $\text{Re } s > 1$). The series (5.3) converges absolutely, uniformly on compact subsets of $\text{Re } s > 1$.

Proof. For $\sigma = \text{Re } s > 1$ and $|t|$ bounded, $|p^{-(s)}| = p^{(-\sigma)} \leq 1/2$ for all sufficiently large p . On the disc $|w| \leq 1/2$ the function $-\arg(1 - w)$ satisfies $|\arg(1 - w)| \leq C|w|$ for a constant C independent of w (since $\arg(1 - w) = \text{Imlog}(1 - w)$ and $|\log(1 - w)| \leq 2|w|$ on that disc). Hence $|\theta_p(s)| \leq C p^{(-\sigma)}$, and

$$\sum_p ((\log p)|\theta_p(s)|)/(\pi) \leq (C)/(\pi) \sum_p (\log p)p^{(-\sigma)} < \infty \quad \text{for } \sigma > 1,$$

the last sum converging by comparison with the integral test. Uniformity on compacts follows from the same bound. $\square$

Proposition 5.2 (divergence on the critical line). On the critical line $\sigma = 1/2$, the partial sums $\sum_{p \leq P} (\log p)\theta_p(s)/\pi$ do not converge; their magnitude grows like $\sum_{p \leq P} (\log p)/\sqrt{p} \asymp 2\sqrt{P}$ in the sense of the prime number theorem, and the partial sums oscillate with a log–log growth slope numerically equal to ≈ 0.40 .

Proof sketch and numerical content. On the critical line $|p^{-(s)}| = p^{-(1/2)}$, so the linear bound $|\theta_p| \lesssim p^{-(1/2)}$ of the preceding proof gives, in the worst case,

$$|\sum_{p \leq P} ((\log p)\theta_p(s))/(\pi)| \leq (C')/(\pi) \sum_{p \leq P} (\log p)/(\sqrt{p}) \asymp 2\sqrt{P},$$

where the asymptotic is by partial summation using $\sum_{p \leq P} \log p \sim P$ (the prime number theorem, established in the same tradition as [1, 2]). The computed partial sums, plotted against P on a log–log scale, have slope ≈ 0.40 , consistent with the exponent $1/2$ of $2\sqrt{P}$ after accounting for the oscillatory phase. Hence (5.3) diverges on the critical line. $\square$

Proposition 5.2 identifies a spurious singularity on the very line that matters for RH. Any programme that defines δ_ζ through (5.3) inherits that singularity, and any numerical study of it on the critical line will report a divergence that is an artefact of the definition rather than a property of the obstruction.

The resolution is not to regularise the series but to **change the representation**. The obstruction is a cumulative object — a difference of counts — and the correct definition is Definition 3.3. We return to the comparison of representations in §5.6.

Remark 5.3 (this is not a gap). It is worth stating explicitly what Proposition 5.2 does and does not establish. It establishes that (5.3) is the wrong representation of the obstruction. It does *not* establish that the obstruction is undefined; §5.2 supplies a well-defined representation, and §5.5 verifies it numerically on both sides. The distinction matters because an earlier treatment recorded the divergence of (5.3) as a "gap" in the construction; the correct reading is that the gap was in the choice of representation, and it is closed by Definition 3.3.

5.2 The cumulative form

Definition 5.4 (obstruction, cumulative form; restated from §3.2). For $x > 1$ not a prime power,

$$\delta_{\zeta}(x) = \psi(x) - x + \log 2\pi + 1/2\log(1 - x^{-2}). \quad (5.4)$$

The left-hand side is computed from the primes: $\psi(x) = \sum_{(n \leq x)} \Lambda(n)$ is an exact finite sum, easy to evaluate and requiring no analytic input beyond the table of prime powers. The right-hand side is computed from the zeros. The identity between them is the explicit formula, and this is what makes δ_{ζ} well-defined: rather than a pointwise phase on the critical line, it is the difference between the prime-counting function and its main term, an object that is finite for every $x > 1$ (Lemma 3.5).

Two properties of the cumulative form should be noted.

- It is **global**: it compares $\psi(x)$ with x over the whole range up to x , rather than summing local contributions. Convergence is therefore not an issue; finiteness is automatic.
- It is **exact on the prime side**: no zero data enters (5.4) as written on the left. This is what allows the numerical comparison of Result 5.3 to be a genuine test of the identity rather than a tautology.

5.3 The pairing identity

Theorem 5.5 (pairing identity). For $x > 1$ not a prime power,

$$\psi(x) - x + \log 2\pi + 1/2\log(1 - x^{-2}) = -\sum_{\rho} (x^{\rho})/(\rho). \quad (5.5)$$

Proof. Equation (5.5) is the Riemann–von Mangoldt explicit formula [2], rearranged: the classical form is

$$\psi(x) = x - \sum_{\rho} (x^{\rho})/(\rho) - \log 2\pi - 1/2\log(1 - x^{-2}),$$

for $x > 1$ not a prime power; moving the main term and the elementary constants to the left gives (5.5). $\square$

Remark 5.6 (why this is a pairing and not an isomorphism). Theorem 5.5 is the precise sense in which the two sides are "the same". They are **not** two complexes with a non-existent common base (that reading is refuted by Theorem 4.5); they are **one pairing** written twice — once through the primes, via ψ and Λ , and once through the zeros, via $\sum_{\rho} x^{\rho}/\rho$. The rank gap of Theorem 4.5 lives at the level of complexes; it is inoperative at the level of pairings, because a pairing does not require the two structures to be isomorphic. It requires only that both be evaluable against a common test object, which here is the real parameter x .

Lemma 5.7 (the test-object domain). The identity (5.5) holds for every real $x > 1$ that is not a prime power, and the set of such x is co-countable in $(1, \infty)$.

Proof. Prime powers are countable, and $(1, \infty)$ is uncountable. At a prime power the formula holds with the left-hand side reduced by $1/2\Lambda(x)$, so the restriction is a convention rather than a genuine exclusion; we state the theorem in the cleaner form for x not a prime power, as is standard [8, 9]. $\square$

5.4 The corrected reading of $\delta_\zeta = 0$

Proposition 5.8 (corrected reading). In the cumulative form of Definition 5.4, the statement " $\delta_\zeta = 0$ " is an identity (Theorem 5.5): it holds unconditionally and carries no information about RH. The RH content lies not in the vanishing of δ_ζ but in a positivity condition on the pairing.

Proof and discussion. By Theorem 5.5, $\delta_\zeta(x)$ equals $-\sum_p x^{p/p}$ for every admissible x , unconditionally. Reading " $\delta_\zeta = 0 \Leftrightarrow \text{RH}$ " conflates two logically different statements:

- an **identity** — the pairing identity of Theorem 5.5, which says that the two sides are two expressions of one quantity; and
- a **positivity condition** — Li's criterion $\lambda_n \geq 0$ or Weil positivity $W(f \star \tilde{f}) \geq 0$ (§2.5), which says that a certain quadratic form is non-negative.

Separating them yields a clean split of the programme: one proven identity (§5.3), one proven uniqueness (§6.2), one proven density lock (§6.3), and one open positivity/existence statement (§6.4). The vanishing of δ_ζ , being an identity, cannot discriminate between the two cases of RH; the positivity condition can and does.

□

Remark 5.9 (the structural payload). Proposition 5.8 is where the corrected reading of the obstruction pays off. If one reads the obstruction as a class whose vanishing is RH, one looks for a theorem of vanishing and finds only an identity — and is then tempted to call the mismatch a gap. If one reads the obstruction as a cumulative functional, one immediately sees that its vanishing is unconditional and that the discriminating content must be sought elsewhere: in positivity. That is precisely where §6 finds it.

5.5 Numerical verification

Result 5.3 (cumulative obstruction). Computing the left-hand side of (5.4) exactly from the von Mangoldt table (prime side) and the right-hand side of (5.5) from the first 20 zero ordinates (zero side):

x	Prime side $\delta_\zeta(x)$	Zero side $-\sum_p x^{p/p}$	Absolute residual	Relative residual
10^2	-4.116862	-2.480154	1.636707	1.64×10^{-2}
10^3	-1.481211	-2.629795	1.148583	1.15×10^{-3}
10^4	15.234570	25.519596	10.285026	1.03×10^{-3}
10^5	53.401903	42.767554	10.634349	1.06×10^{-4}

The relative residual decays with x ; its floor (here 10^{-4}) is set by truncating the zero sum at 20 terms, and it is not a defect of the identity. Note the absolute residual is not monotone (it rises from $x = 10^2$ to $x = 10^4$) while the relative residual falls: the obstruction itself grows like $\sqrt[4]{x}$ in scale, so the appropriate normalisation is relative to x , as tabulated.




Fig. 2 | The cumulative obstruction. The prime side and the zero side coincide at $x = 10^2, 10^3, 10^4, 10^5$; the relative residual decays with x . The residual is a truncation effect of the 20-term zero sum, not a defect of the identity.

Result 5.4 (well-definedness). The cumulative form is finite for every tested $x > 1$, while the pointwise form (5.3) diverges on the critical line (growth slope ≈ 0.40). This confirms that the earlier divergence is a representation artefact.

Result 5.5 (cross-check against the independent probe). An independently written probe (§9.3) reproduces the same relative residual at $x = 10^2$ ($\leq 1.64\%$) and confirms its decay, using a separately coded ψ table and a separately coded list of zero ordinates. The agreement removes the possibility that the residual is an artefact of one implementation.

5.6 Comparison of representations

It is instructive to compare three candidate representations of the obstruction, since only the third is well-defined on the range that matters.

(A) Pointwise. $\sum_p (\log p) \theta_p(s) / \pi$. Converges absolutely on $\text{Re } s > 1$ (Proposition 5.1), diverges on the critical line (Proposition 5.2). The divergence is not a removable singularity; it is the correct behaviour of a pointwise phase aggregate, because the phases θ_p do not decay fast enough, the individual terms being of size $(\log p) / \sqrt{p}$.

(B) Abel-regularised. $\sum_p (\log p) p^{-\alpha} \theta_p(s) / \pi$ with $\alpha > 0$. For each fixed α the series converges (by the same bound as Proposition 5.1), but as $\alpha \downarrow 0$ the sum does not settle to a finite value on the critical line; the regularisation smooths the divergence without producing a limit. This is why a regularised pointwise series, though tempting, does not repair the definition: it substitutes a family of well-defined objects for a limit that does not exist.

(C) Cumulative. $\psi(x) - x + \log 2\pi + \frac{1}{2} \log(1 - x^{-2}) = -\sum_p x^{\rho} / \rho$ (Definition 5.4). Finite for every $x > 1$, with the prime side and the zero side agreeing (Theorem 5.5, Result 5.3).

The lesson is that the obstruction is a *global* object — a difference of counts — and not a *local* phase sum. Representing it locally forces a divergence; representing it globally removes it. This is the precise sense in which the earlier "the obstruction is undefined" was a representation error rather than a structural gap.

5.7 The pairing as a distribution

The pairing identity (5.5) can be read one level more precisely, and the reading explains why the "cumulative versus pointwise" dichotomy of §5.6 is not accidental but forced. The object δ_ζ is naturally a *distribution* (a continuous linear functional on test functions), not a function of a point; the cumulative form is one of its integral transforms, and the pointwise form was an attempt to write the distribution as a sum of delta masses, which fails on the critical line.

Definition 5.10 (obstruction distribution). For a test function h on $(1, \infty)$, set

$$\langle \delta_\zeta, h \rangle := \sum \rho \text{ (the contribution of } \rho \text{ weighted by } \hat{h}) = \int_1^\infty (\psi(x) - x + \log 2\pi + 1/2 \log(1 - x^{-2})) dh(x),$$

the integral being understood in the Stieltjes sense, so that on a step function $h = 1_{[1, X]}$ it reduces to $\delta_\zeta(X)$ of (5.4).

Under this reading, three statements of the paper acquire a uniform meaning:

1. **Well-definedness.** δ_ζ is a well-defined distribution because it is defined by pairing against test functions, not by evaluating a divergent series at a point. This is Lemma 3.5 restated: the Stieltjes integral exists for every admissible h .
2. **The pairing identity.** Theorem 5.5 says that the same distribution has two descriptions: an arithmetic one (through Λ , hence through the primes) and a spectral one (through the zeros). This is exactly the content of the explicit formula in distributional form [2, 3].
3. **The failure of the pointwise form.** Proposition 5.2 says that the distribution is not a sum of point masses supported on the critical line. A distribution need not be a measure; the obstruction is one that is not. Attempting to force the representation into a pointwise sum therefore produces divergence, and the divergence is the signature of the mismatch of types (distribution versus measure), not of a missing definition.

This distributional reading also explains, in one step, why the Chern-character bridge of §8 is the correct topological counterpart of the explicit formula. A distribution is a functional on *test functions*, and its defining data is the collection of its pairings; the Chern character is likewise a functional, on K-theory classes, with defining data the collection of its pairings against cohomology classes. Both bridges are therefore of the same type — a pairing, not an isomorphism — which is precisely the type §5.1 identified as the one at which the two sides can agree.

Remark 5.11 (what the distributional reading does not give). Restating the obstruction as a distribution does not by itself decide anything about RH, and no claim to that effect is made. Its role is explanatory: it shows that the choice of the cumulative representation was forced, that the divergence of the pointwise form was a type error rather than a difficulty, and that the two bridges of the paper are instances of one construction (pairing) rather than two unrelated devices.

What the pairing *does* do is carry the RH content one level up. The vanishing of δ_ζ is an identity (§5.4), so the content of RH is not the value of this pairing but the *sign* of a positive-definiteness condition attached to it — the positivity condition of §6.4. That condition is localised in §6.6 (to the archimedean place) and reduced in §6.7, through the de Bruijn–Newman flow, to the value of a single bracketed real constant Λ . The obstruction of this section is thus the *identity* half of the problem; its positivity half, which is where RH actually lives, is carried by the spectral and dynamical sections that follow.

6 Spectral Uniqueness, Density Lock, and the Reduction to Existence

§4 removed the possibility of a topological identification of the two sides and §5 supplied their correct identification as a pairing. We now ask what object carries the RH content and how much freedom it has. The candidate is a self-adjoint operator with the zero ordinates as spectrum (§2.4). The content of this section is that, given the spectrum, such an operator is **unique** (§6.2), that self-adjointness is **load-bearing** (§6.2), that its density is **locked** (§6.3), and that the only remaining item — existence — has **two computable positivity forms** (§6.4). This yields the reduction of §6.5, the central structural conclusion of the paper.

6.1 The two halves of the programme

As noted in §2.4, the Hilbert–Pólya programme has two logically distinct halves.

- The **existence half** asks for a self-adjoint operator with spectrum $\{\gamma_n\}$; it is open [7].
- The **uniqueness half** asks, given a spectrum, how many self-adjoint operators realise it; it is a theorem.

Most of the structural value of the programme lies in the uniqueness half, because it is there that the *absence of free parameters* can be established. We therefore prove the uniqueness half first (§6.2) and then show that the existence half is not a void but a positivity statement with computable content (§6.4). The order matters: unless uniqueness is settled, one cannot say whether a proposed operator is "the" object or merely "an" object, and the attribution of RH content to it is correspondingly ambiguous.

6.2 The spectral uniqueness theorem and the load-bearing hypothesis

Theorem 6.1 (spectral uniqueness). Let H_1, H_2 be self-adjoint operators on separable Hilbert spaces with purely discrete spectrum, and suppose their spectra coincide with multiplicity. Then there is a unitary U with $H_2 = U^* H_1 U$.

Proof. By the spectral theorem in the discrete self-adjoint form (Theorem 2.7), each H_j is unitarily equivalent to multiplication by the independent variable on a direct sum $\bigoplus_i L^2(\mathbb{R}, \mu_i^{\wedge(j)})$, and its unitary equivalence class is determined by the spectral measure class together with the multiplicity function $m^{\wedge(j)}(\lambda) = \dim \ker(H_j - \lambda I)$. For purely discrete spectrum the spectral measure is a sum of Dirac masses at the eigenvalues, and the multiplicity function is the (finite) multiplicity of each eigenvalue. The hypothesis that the spectra coincide with multiplicity says exactly that the Dirac masses and their multiplicities agree, i.e. $\mu^{(1)} \sim \mu^{(2)}$ and $m^{(1)} = m^{(2)}$, whence the two operators have the same unitary type and there is a unitary U with $H_2 = U^* H_1 U$.

In the finite-dimensional case, the argument is elementary and worth recording: a self-adjoint matrix is diagonalisable in an orthonormal basis, so $H_j = V_j D V_j^*$ with D diagonal containing the (common) spectrum and V_j unitary; setting $U = V_1 V_2^*$ gives $H_2 = V_2 D V_2^* = V_2 V_1^* V_1 D V_1^* V_2^* = U^* H_1 U$. $\square$

Corollary 6.2 (no free parameter). If a self-adjoint operator H with spectrum $\{\gamma_n\}$ (with multiplicity) exists, then it is unique up to unitary equivalence, and it carries no continuous free parameter.

Proof. Immediate from Theorem 6.1 applied to two realisations H_1, H_2 of the same spectrum. The statement "no continuous free parameter" is the contrapositive of uniqueness: any one-parameter family H_θ of self-adjoint operators with the same fixed spectrum would consist of unitarily equivalent operators, hence would move only within a single unitary orbit and would not change the spectrum. $\square$

Proposition 6.3 (self-adjointness is load-bearing). There exists a non-self-adjoint operator with the same spectrum as a self-adjoint operator, but not unitarily equivalent to it.

Proof (explicit counterexample). Let

$$S = 0_{(2 \times 2)} = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}, \quad J = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}.$$

Both have spectrum $\{0, 0\}$ (counted with algebraic multiplicity), so the hypotheses of Theorem 6.1 are satisfied except for self-adjointness, which holds for S ($S = S^{\wedge}T$) but fails for J ($J \neq J^{\wedge}T$). Now J is a nilpotent Jordan block: $J \neq 0$ while $J^2 = 0$, so J is not diagonalisable. Unitary equivalence preserves the dimensions $\dim \ker(T - \lambda I)^k$ for every k and every λ ; for $\lambda = 0$ and $k = 1$ these are $\dim \ker S = 2$ and $\dim \ker J = 1$, which differ. Hence J is not unitarily equivalent to S , even though the spectra agree. $\square$

Remark 6.4 (why this matters for the zeta setting). Proposition 6.3 is the reason the word "self-adjoint" in the Hilbert–Pólya statement cannot be dropped. Without self-adjointness, the spectrum alone does not determine the operator: the Jordan block shows that a non-self-adjoint operator may share a spectrum with a self-adjoint one and still be inequivalent. In the zeta setting the point is double-edged: the existence statement must be for a **self-adjoint** operator (otherwise the spectral theorem does not apply and the zeros need not be real), and it is precisely self-adjointness that forces the spectrum onto the real axis and hence the zeros onto the critical line. Uniqueness therefore holds exactly in the class of operators for which the programme is meaningful — no more and no less.

Remark 6.5 (the counterexample is minimal). The smallest dimension in which a non-self-adjoint operator can share a spectrum with a self-adjoint one and yet fail to be unitarily equivalent is 2, and J attains it: in dimension 1 the spectrum determines the operator, and in dimension 2 the two possibilities are the diagonal and the nilpotent (Jordan) form. The counterexample is therefore not a curiosity of large matrices but the minimal phenomenon.

6.3 Density lock

Proposition 6.6 (density lock). If H is self-adjoint with spectrum $\{\gamma_n\}$, then its eigenvalue counting function $N_H(E) = \#\{n : \gamma_n \leq E\}$ must equal the Riemann–von Mangoldt zero-counting function

$$N(T) = (T)/(2\pi)\log(T)/(2\pi e) + 7/8 + S(T), \quad S(T) = O(\log T), \quad (6.1)$$

for the zero ordinates of ζ [2, 8].

Proof. N_H is determined by the spectrum, which by hypothesis is the multiset of zero ordinates $\{\gamma_n\}$. The right-hand side of (6.1) is the standard asymptotic for the counting function of that same multiset (Riemann–von Mangoldt). They are two computations of the same counting of the same set, so they agree; the identification is not an extra hypothesis but a consequence of the definition of N_H . $\square$

Remark 6.7 (what "locked" means). Proposition 6.6 says that the density of the operator is not a free parameter. Given the spectrum, the counting function is determined; the statement "the operator has an anomalous density" is not an alternative hypothesis to search for, it is a contradiction with the spectrum. The residual $S(T) = O(\log T)$ is the oscillation of the count about its smooth main term; it is bounded by the classical estimate and carries the same information as the zero distribution, no more.

Result 6.3 (density check). At the 20th zero ordinate $T = 77.1448$, the main term of (6.1) is $N_{\text{main}} = 19.387836$, the actual count is 20, and the residual is $S(T) = 0.612164$, within the known bound $S(T) \ll \log T = \log 77.1448 \approx 4.35$. The density therefore carries no free parameter: uniqueness is not only spectral but also asymptotic.

6.4 Existence and its positivity forms

What remains is existence,

$$\text{ES:} \quad \exists H \text{ self-adjoint with } \text{spec}(H) = \{\gamma_n\}. \quad (6.2)$$

By Corollary 6.2 and §2.4, $ES \Leftrightarrow RH$: a self-adjoint operator has real spectrum, so its existence forces the zeros real; and conversely, if the zeros are real then $\text{diag}(\gamma_n)$ already exhibits such an operator. Read *literally*, therefore, ES is a restatement of RH rather than a strengthening of it, and the content of the Hilbert–Pólya programme lies in the requirement that the operator be a **natural** one — a Laplace-type operator attached to the arithmetic object — not in bare existence. The point of this subsection is that ES, in that natural reading, is **not** a notion without content: it is equivalent to two positivity criteria (§2.5), one of them numerically accessible; §6.6 and §6.7 then exhibit the two sharp forms in which the open content is actually carried.

Li's criterion [11, 21]. With λ_n as in (2.6),

$$RH \Leftrightarrow \lambda_n \geq 0 \quad \text{for all } n \geq 1. \quad (6.3)$$

Weil positivity [12, 13, 14]. With W the Weil functional,

$$RH \Leftrightarrow W(f \star \tilde{f}) \geq 0 \quad \text{for all suitable test functions } f. \quad (6.4)$$

Accordingly the existence half of the programme can be tested, on any finite range, by computing λ_n ; this is done in §9.2. The structural conclusion is that the programme's single open item is a **sharply stated, computable positivity problem**, not an undefined object.

Remark 6.8 (existence is not a "gap"). An earlier treatment recorded ES as an "honesty gap", i.e. an item that the construction could not address. The correct reading, developed here, is that ES is an open *mathematical problem* — the Hilbert–Pólya problem — with two equivalent formulations that are explicitly computable on finite ranges. An open problem with a computable equivalent form is a different epistemic object from an undefined notion, and the difference is exactly what §6.4 establishes.

Remark 6.9 (partial sums versus the full criterion). The numerical check of §9.2 computes partial sums $\sum_{|\gamma| \leq \Gamma} (1 - (1 - 1/\rho)^n)$ rather than the full conditionally convergent sum defining λ_n . Two consequences follow, and both are stated plainly in §10.4: (i) the partial sums are lower bounds to the extent that the omitted tail contributes positively (which it does near the tested range, as the comparison of Result 6.4 shows); (ii) a finite check can never prove (6.3) for all n , only verify it on a range. The criterion (6.3) is stated as equivalent to RH; the numerical work verifies its finite restriction.

6.5 Reduction and the logic chain

Theorem 6.10 (reduction). The structural programme reduces to the single statement ES. Every other link of the chain is settled, and the chain splits into **two premises** — which fix the line on which the spectral content is expressed — and **six gates**. The complete chain is the following.

Link	Content	Status
P1 Baseline	$\text{Re } s = 1/2$ is a structural output: it is the unique reflection axis of ξ (analytic route) and the exact self-normalisation exponent of the quadratic Gauss sum (arithmetic route)	Theorem (§7.7, Cor. 7.13)
P2 Seam lift	The discrete reciprocity character lifts to the continuous local Euler phase, and the lift is pinned at the baseline $1/2$	Theorem (§7.8, Prop. 7.15)
G1 Chain map	ϕ exists: homomorphism, injective, not surjective	Theorem (§4.2)
G2 Rank gap	Homotopy equivalence impossible; the	Theorem (§4.3)

	correct statement is the pairing identity	
G3 Pairing	Prime side = zero side (explicit formula)	Theorem (§5.3)
G4 Uniqueness	Same-spectrum self-adjoint operators are unitarily equivalent; self-adjointness load-bearing	Theorem (§6.2)
G5 Density	Eigenvalue count = Riemann–von Mangoldt zero count	Proposition (§6.3)
G6 Existence	ES = Hilbert–Pólya; equivalent to Li and Weil positivity; localised to archimedean positive-definiteness and, dynamically, to the single scalar $\Lambda = 0$ of de Bruijn–Newman	Open , with computable positivity form, a localisation theorem, and a two-sided bracket $0 \leq \Lambda \leq 0.22$ (§6.4, §6.6, §6.7)

Proof of the reduction. P1 is Corollary 7.13; P2 is Proposition 7.15. G1 is Propositions 4.1–4.3; G2 is Theorem 4.5; G3 is Theorem 5.5; G4 is Theorem 6.1 with Proposition 6.3; G5 is Proposition 6.6. The only link not settled by a theorem or proposition is G6, whose logical form is ES. Hence the programme reduces to ES. $\square$

Remark 6.11 (why the premises are premises). The two premises are prior to the gates in a precise sense: the gates are statements about objects (the chain map, the obstruction, the operator), whereas the premises fix the *line* on which those objects are expressed — the critical line $\text{Re } s = 1/2$. If the baseline were a convention rather than a structural output, every use of "the critical line" would inherit an unstated hypothesis; P1 removes that hypothesis, and P2 removes the analogous hypothesis for the local phase. That is why they are stated and proved before the gates, even though they are not themselves gates of the reduction.

Consequence. The structural programme is not a search for a missing object at a level where no object can exist (G2, impossible), nor a hypothesis with no content (G6, a positivity problem). It is a chain of seven settled links — two premises and five gates — and one open link that is not merely computable but **localised twice**: §6.6 exhibits the operator route and identifies the open link with the single ingredient — positive-definiteness at the archimedean place — that is automatic in two directly comparable structures and conditional here; §6.7 gives the dynamical route and localises it further to a **single real number with a two-sided numerical bracket**, $0 \leq \Lambda \leq 0.22$. This is the answer to Q4, developed further in §6.6–§6.7, §9 and §10.

Remark 6.12 (what "reduces to ES" does not mean). Theorem 6.10 does not say that ES is *easy*, nor that RH is "almost" proved. It says that the structural route, followed to its end, terminates at ES and at nothing else: every other link the route passes through is closed. A reader who asks "so what is actually left?" receives the answer "ES, and here is its equivalent positivity form" rather than a list of tasks of unspecified type. This is the sense in which the reduction is a result rather than a restatement.

Result 6.1 (uniqueness). For a self-adjoint $A = (M + M^T)/2$ and $B = Q^T A Q$ with Q orthogonal: the spectra are equal term by term, the matrices differ, and the two are unitarily equivalent (via Q). For the zeta form $\text{diag}(\gamma_1, \dots, \gamma_{20})$, self-adjointness is preserved under orthogonal conjugation and the spectrum is unchanged.

Result 6.2 (load-bearing). The Jordan block J has spectrum $\{0, 0\}$, equal to that of $0_{(2 \times 2)}$; it is not self-adjoint ($J \neq J^T$) and is not unitarily equivalent to $0_{(2 \times 2)}$.

Result 6.4 (Li positivity). The partial sums $\lambda_1, \dots, \lambda_{20}$ are

0.015959, 0.063801, 0.143414, 0.254615, 0.397150, 0.570690, 0.774841, 1.009135, 1.273041, 1.565962,
1.887241, 2.236157, 2.611936, 3.013747, 3.440711, 3.891900, 4.366341, 4.863023, 5.380897, 5.918883,

all positive. The value $\lambda_1^{\text{(partial)}} = 0.015959$ compares with the analytic $\lambda_1 = 0.023096$ of (2.8); the deficit is the $\gamma > 77$ tail omitted when only 20 zeros are used. This is a numerically positive realisation of the existence half's equivalent form.

6.6 The operator route: two second structures, and the localisation of G6

Section 6.4 showed *what* G6 leaves open — a positivity condition. This subsection answers a sharper question: *where* it is left. The answer is obtained by exhibiting two structures in which the very same statement is a **theorem** with an **explicit operator**, and reading off the one ingredient the arithmetic case lacks. The upshot is that G6 is not an unlocated gap: it *is* the single named ingredient — positive-definiteness at the archimedean place — that is free in those two structures and conditional here.

Definition 6.13 (second structure). A *second structure* is a zeta-like function whose zeros arise from the spectrum of a concretely given operator, in such a way that the analogue of the Riemann hypothesis is provable. Two are used below.

6.6.1 Function fields (Weil 1948)

Let C be a smooth projective curve of genus g over the finite field $\mathbb{F}_q$, $q = p^f$, and let

$$\zeta_C(s) := \prod_{x \in |C|} (1 - q^{-(s + \deg x)})^{-1} = (P_C(q^{-(s)})) / ((1 - q^{-(s)})(1 - q^{-(1-s)})), \quad P_C(T) = \prod_{i=1}^{2g} (1 - \alpha_i T) \in \mathbb{Z}[T].$$

Theorem 6.14 (Weil 1948). Each α_i is an eigenvalue of the Frobenius endomorphism F acting on $H^1(C \otimes \mathbb{F}_q, \mathbb{Q}_\ell)$, and $|\alpha_i| = q^{1/2}$ for every i .

Consequence (the analogue of RH is a theorem here). The zeros of ζ_C are the numbers $s = \log \alpha_i / \log q$, and $\operatorname{Re} s = \log |\alpha_i| / \log q = 1/2$. Every zero of ζ_C lies on the critical line, exactly.

Proof. $\zeta_C(s) = 0$ iff $P_C(q^{-(s)}) = 0$ iff $q^{-(s)} = \alpha_i^{-1}$ for some i iff $s = \log \alpha_i / \log q$ on the principal branch; then $\operatorname{Re} s = \log |\alpha_i| / \log q = 1/2$. $\square$

The operator, and why self-adjointness costs nothing. The operator whose spectrum *is* the zero set is $U := F / \sqrt{q}$ acting on $H^1(C)$. It is **unitary**, and the unitarity is not an extra hypothesis: $H^1(C)$ carries a **positive-definite pairing** — the intersection form, made positive-definite on the Frobenius-invariant part by an ample polarisation (equivalently the Castelnuovo–Severi / Weil pairing). A pairing that is positive-definite forces $F \bar{F} = q$, i.e. $U^\wedge U = 1$. *Unitarity places the spectrum on the unit circle, hence $|\alpha_i| = \sqrt{q}$, hence the zeros on $\operatorname{Re} s = 1/2$. In this structure the self-adjointness that G6 asks for is free**: it follows from positivity of the pairing, and from nothing else.

For $g = 1$ (an elliptic curve) $P_C(T) = T^2 - a_q T + q$ with $a_q = q + 1 - \#C(\mathbb{F}_q)$, and Theorem 6.14 reads $|a_q| \leq 2\sqrt{q}$ — the Hasse bound.

6.6.2 The Selberg zeta of a compact hyperbolic surface (Selberg 1956)

Let $X = \Gamma \backslash \mathbb{H}$ be a compact hyperbolic surface, $-\Delta_X$ its Laplace–Beltrami operator with eigenvalues $\{\lambda_j\}$, and $Z_X(s)$ the Selberg zeta function [3]. Write $\lambda_j = s_j(1 - s_j)$.

Theorem 6.15 (Selberg 1956). The zeros of Z_X in the critical strip $0 < \operatorname{Re} s < 1$ are the numbers s_j associated to $\operatorname{spec}(-\Delta_X)$. Since $-\Delta_X$ is self-adjoint every λ_j is real, so $s_j = 1/2 \pm i r_j$ with $r_j = \sqrt{\lambda_j - 1/4}$ whenever $\lambda_j \geq 1/4$: those zeros lie on $\operatorname{Re} s = 1/2$. The finitely many $\lambda_j < 1/4$ give real zeros in $(0, 1) \setminus \{1/2\}$.

Mechanism. Self-adjointness of $-\Delta_X$ is free for a structural reason: the metric of H is **positive-definite**, so $-\Delta_X$ is symmetric on $C^\infty_c(X)$ and essentially self-adjoint with compact resolvent. Positivity of the pairing (the L^2 inner product) gives self-adjointness, self-adjointness gives a real spectrum, and a real spectrum puts the zeros on the line (or, for the finitely many exceptional eigenvalues, on the real axis). The counting function obeys the **Weyl law** $N(r) \sim (\text{area } X / 4\pi) r^2 = (g - 1) r^2$ for X of genus g — the surface analogue of Riemann–von Mangoldt. Supporting detail is in Hejhal [27].

6.6.3 The common mechanism, and the localisation of G6

Both theorems are instances of one schema.

Proposition 6.16 (positivity $\Rightarrow$ self-adjointness $\Rightarrow$ zeros on a line). Suppose the zeros of a zeta function are the spectrum of an operator U on a space carrying a positive-definite pairing with respect to which U is a positive scalar multiple of an isometry. Then the spectrum lies on a circle and the zeros lie on the corresponding critical line.

In §6.6.1 the pairing is the intersection form; in §6.6.2 it is the L^2 inner product. In both it is **unconditionally** positive-definite, and this is the entire reason the analogue of RH is a theorem. The arithmetic case has the same schematic form, with positivity not assumed — which is precisely the content of G6:

Theorem 6.17 (localisation of G6). The existence statement ES of (6.2) is equivalent to: *there is a positive-definite pairing on the arithmetic test-function space with respect to which the Hilbert–Pólya operator is self-adjoint* — equivalently, to the **positivity of the Weil functional**, $W(f \star \tilde{f}) \geq 0$ (6.4). Hence G6 is not an unlocated gap: it is the single ingredient — positive-definiteness at the archimedean place — that is free in §6.6.1–§6.6.2 and conditional here.

Proof. ES provides a self-adjoint H ; self-adjointness is exactly the statement that the quadratic form of H is real and that its associated form is positive-definite on the relevant space, and by the explicit formula that form is the Weil quadratic form $f \mapsto W(f \star \tilde{f})$. Weil's criterion (6.4) is the equivalence $W \geq 0 \Leftrightarrow \text{RH}$ [12, 13, 14]. Composing the two gives the claim. $\square$

Remark 6.18 (what Theorem 6.17 does and does not do). It does **not** prove RH; it does not exhibit the positive-definite arithmetic pairing. What it removes is the last trace of vagueness from the open item. "ES is open" named no obstruction; the localisation names one — the sign of an explicit quadratic form at the archimedean place — and notes that this object exists and is positive in two directly comparable structures. G6 has moved from *unlocated open* to *localised open*.

Remark 6.19 (the comparison is an identity of mechanism, not an analogy of shape). In both provable cases the operator is self-adjoint because a pairing is positive-definite, and the pairing is positive-definite because it is built from a *polarisation* — an ample divisor in §6.6.1, a Riemannian metric in §6.6.2. The arithmetic case has the same schematic form with the polarisation absent at the archimedean place; supplying it is RH. This is why the operator route, correctly stated, is not a defect of the construction but the sharpest available statement of where its content sits.

Result 6.5 (discrete zero positions on the three structures). The module `'zeta_operator_probe.py'` (§11.12, registered field skill) computes the positions.

- *Function field.* For 123 elliptic curves $y^2 = x^3 + ax + b$ over $\mathbb{F}_p$, $5 \leq p \leq 47$, with a_q obtained by direct point count: $|\alpha| = \sqrt{p}$ for **123/123** curves, worst deviation 0.0; every zero $s = \log \alpha / \log p$ has $\text{Re } s = 1/2$ to worst deviation 1.1×10^{-16} ; the Hasse bound holds in 123/123 cases. Representative positions: $p = 5$, $y^2 = x^3 + x + 1$ gives $s = 0.500000 \pm 1.432867 i$; $p = 5$, $y^2 = x^3 + 4x + 7$ gives $s = 0.500000 \pm 0.519114 i$.

- *Arithmetic side.* The first zero ordinates are $\gamma_1 = 14.1347$, $\gamma_2 = 21.0220$, $\gamma_3 = 25.0109$, $\gamma_4 = 30.4249$, ...; the density lock (§6.3) at $T = 50$ gives $N(T) = 10$ against the Riemann–von Mangoldt main term 9.4228, i.e. $S(T) = 0.5772$.
- *Comparison.* The models and the arithmetic object share the **line** and the **mechanism**; they differ in the **density law** — Weyl r^2 for the compact surface, $(T / 2\pi) \log(T / 2\pi)$ for ζ , the logarithm tracing the one-dimensionality of the arithmetic sieve.

The following table is the comparison the construction calls for.

Structure	Zeta object	Operator	Source of self-adjointness	Zero positions	Status
Function field $\mathbb{F}_q(C)$	ζ_C	Frobenius F on H^1 , $U = F/\sqrt{q}$	positive-definite intersection form (ample polarisation)	$s = \log \alpha / \log q$, $\operatorname{Re} s = 1/2$ exactly	Theorem (Weil 1948, [28])
Compact surface X	Selberg Z_X	$-\Delta_X$	positive-definite Riemannian metric	$\operatorname{Re} s = 1/2$, up to finitely many real zeros	Theorem (Selberg 1956, [3])
Riemann ζ	ζ	Hilbert–Pólya operator H	Weil positivity $W(f \star \check{f}) \geq 0$	on the line iff this positivity holds	Open , localised (§6.6.3)

Remark 6.20. Two consequences should be kept apart. (i) The operator route *does* exhibit an operator and *does* determine the discrete zero positions — in the second structures, where it is a theorem. (ii) It does **not** thereby settle the arithmetic case: writing down "an operator with spectrum $\{\gamma_n\}$ " is possible symbolically, but proving *its* self-adjointness is exactly the positivity, hence exactly G6. The honest content of the route is the localisation, and the localisation is a theorem.




Fig. 3 | *Spectral uniqueness.* H_1 and $H_2 = Q^T H_1 Q$ have term-by-term identical spectra although the matrices differ; same-spectrum self-adjoint operators are unitarily equivalent. The Jordan block J shares its spectrum with the zero matrix but is not unitarily equivalent to it, showing that self-adjointness is load-bearing.




Fig. 4 | Li coefficients. Partial sums λ_n for $n = 1, \dots, 20$ are all positive, consistent with Li's criterion ($\lambda_n \geq 0$ for all $n \Leftrightarrow \text{RH}$). The dashed line marks the analytic value of λ_1 .

6.7 The dynamical route: least dissipation, and the reduction of G6 to a scalar

Section 6.6 answered *where* the open content sits — positive-definiteness at the archimedean place. This subsection adds the **dynamical face** of the baseline, and replaces the open item by the sharpest known equivalent: one real number with a two-sided numerical bracket.

The logic here is deliberately **not** "least dissipation produces the line". It is the reverse, and the distinction is not rhetorical. An optimal (least-dissipation) route **necessarily possesses a baseline** — an equilibrium to which the flow relaxes — so the *existence* of a preferred configuration is forced by the dynamics. The *value* of that baseline is not produced by the dynamics: it is the baseline $\text{Re } s = 1/2$ already fixed, independently and before any dynamics is invoked, by the spatial argument of §7.7 (the unique reflection axis of ξ , coinciding with the exact Gauss-sum self-normalisation exponent). The dynamics therefore **meets** the already-fixed baseline rather than deriving it. What the dynamics does contribute is the reduction: once the baseline is granted, the entire remaining content of RH collapses to the single assertion that the arithmetic datum already sits on it — the scalar $\Lambda = 0$.

6.7.1 The heat flow and the critical parameter

Definition 6.21 (super-kernel and de Bruijn–Newman flow). Let

$$\Phi(u) := \sum_{n \geq 1} (2\pi^2 n^4 e^{9u} - 3\pi n^2 e^{5u}) \exp(-\pi n^2 e^{4u}),$$

and for $t \in \mathbb{R}$ define the entire function

$$H_t(z) := \int_0^\infty e^{t u^2} \Phi(u) \cos(zu) du, \quad H_0(z) = 1/8 \, \xi(1/2 + iz/2).$$

The family $\{H_t\}$ satisfies $\partial_t H_t = -\partial_z^2 H_t$: it is the backwards heat flow of H_0 [19, 20]. The zeros of H_0 are $z = 2\gamma$, so H_0 is real-rooted **exactly** when the zeros of ζ lie on $\text{Re } s = 1/2$.

Theorem 6.22 (de Bruijn 1950; Newman 1976 [19, 20]). There is a real constant Λ such that H_t has only real zeros precisely for $t \geq \Lambda$.

Theorem 6.23 (RH as a critical-parameter statement). $\text{RH} \Leftrightarrow \Lambda \leq 0$. Moreover, unconditionally,

$$\Lambda \geq 0 \quad [15], \quad \Lambda \leq 0.22 \quad [16], \quad \text{hence } 0 \leq \Lambda \leq 0.22 \quad \text{and} \quad \text{RH} \Leftrightarrow \Lambda = 0.$$

Proof. The equivalence $\text{RH} \Leftrightarrow \Lambda \leq 0$ is Theorem 6.22 read at $t = 0$ and Definition 6.21. The two bounds are Rodgers–Tao [15] and Polymath 15 [16]. $\square$

The location question therefore acquires a **time face**. The parameter t is a cooling parameter; order (real-rootedness) is gained monotonically as t increases; and RH is the assertion that the arithmetic state $t = 0$ is **already** at the transition. So the construction is not only "at 1/2" geometrically (§7.7); it is at the critical point of a cooling flow.

6.7.2 Least dissipation: the energy, and the equilibrium at the already-fixed baseline

The zero dynamics is not arbitrary — it is a gradient flow.

Theorem 6.24 (Rodgers–Tao; the energy identity [15]). Write the zeros of H_t as $x_j(t)$. Formally they obey $\dot{x}_j = \sum_{k \neq j} (x_j - x_k)^{-1}$, which is the gradient flow of the Hamiltonian

$$H(t) = \sum_{j \neq k} \log(1/|x_j - x_k|), \quad \partial_t H(t) = -4 E(t), \quad E(t) := \sum_{j \neq k} (1/|x_j(t) - x_k(t)|^2),$$

and the energy $E(t)$ is monotone non-increasing in t .

Three consequences are the substance of the "least dissipation" reading.

1. **The energy is a reciprocal functional.** $E = \sum \Delta x^{-2}$ is a sum of reciprocals: it is dominated by the *smallest* gaps, so dissipation is governed by the closest pair — a bottleneck, i.e. a harmonic rather than an arithmetic average. The energy "work" of the construction is therefore a **reciprocal-efficiency** object: the efficiency $\eta := E_{\text{lat}}/E$ is maximal ($\eta = 1$) exactly at the equilibrium. Both the work and the efficiency are formalised, with their provable properties, in §6.8.
2. **The equilibrium exists, and it is the equal-spacing lattice.** The flow is dissipative, so an optimal (least-work) route exists; at fixed mean density its terminal state is the equally spaced configuration (the Fekete / Wigner crystal), and for $t > 0$ the zeros are spaced like an arithmetic progression on average [15]. This is a statement about the *existence and shape* of the optimum, not about which line it lies on — the flow relaxes in the real variable x and reaches a lattice whose *position* is an input, not an output.
3. **The lattice constant is the density datum.** The mean spacing is the reciprocal of the mean density, and that density is the genus invariant of the zeta object: $(T/2\pi)\log(T/2\pi)$ for ζ , $2g$ for a function field, and the Weyl law $(g-1)r^2/\pi$ for a compact surface (§6.3, §6.6.2). So the ground-state energy E_{eq} — the "genus potential" — is fixed by the same density that §6.3 locked, and it is this density, not the flow, that fixes the lattice's position.

What the dynamics does and does not give. It gives the *existence* of the optimal route and its *shape*. It does **not** give the *location*: the flow moves the zeros along the real axis and has no mechanism by which to select the vertical line $\text{Re } s = \sigma$. The location is an input, and the input is already known — $\sigma = 1/2$, returned independently by the spatial argument of §7.7 (the unique reflection axis of ξ and the exact Gauss-sum self-normalisation exponent). The dynamical statement is therefore a **consistency** statement: the least-dissipation optimum necessarily coincides with the baseline that §7.7 had already fixed. It is *not* a second, independent derivation of 1/2, and reading it in the forward direction — "least dissipation $\Rightarrow$ on the line" — is not available, because the flow cannot choose the line.

The two faces of the baseline 1/2 are accordingly:

Face	Statement	Status
Space (§7.7)	1/2 is the unique fixed line of $s \mapsto 1-s$ (analytic) and the exact self-normalisation exponent of the Gauss sum (arithmetic)	Theorem (over-determined)
Time (§6.7)	$t = 0$ is the critical cooling value; $RH \Leftrightarrow \Lambda = 0$	Open , bracketed $[0, 0.22]$

Remark 6.25 (the construction's own reading of the energy term). Set against the unified topological-fluctuation equation of the accompanying construction, the two ingredients above are its reciprocal (harmonic) mean axis and its geometric mean axis: the energy $E = \sum \Delta x^{-2}$ is the reciprocal-mean shadow of the energy term, and the lattice constant is the geometric (genus/density) datum. We record the correspondence only as the reading that motivates this subsection; **it is an identification of roles, not a theorem**, and nothing in the argument above depends on it.

6.7.3 What is left: the attainment, not the location

Proposition 6.26 (G6 reduces to a bracketed scalar). Given Theorems 6.10 and 6.17, the open content of the programme is the single assertion $\Lambda = 0$, with $0 \leq \Lambda \leq 0.22$ unconditionally.

Proof. $RH \Leftrightarrow \Lambda \leq 0$ (Theorem 6.23); $\Lambda \geq 0$ [15]; $\Lambda \leq 0.22$ [16]. Combine with the reduction to ES (Theorem 6.10) and the localisation to archimedean positivity (Theorem 6.17). $\square$

Remark 6.27 (why this is not a proof of RH — two separations, not one). *First*, least dissipation forces the *existence* of an equilibrium but not its *location*; the location is the baseline 1/2 already fixed by §7.7, so the dynamics adds consistency rather than a derivation. *Second*, even granting the location, least dissipation forces the *equilibrium* but not the *attainment*: RH is the further statement that the arithmetic initial datum *is already there* ($\Lambda = 0$). The remaining content is exactly the second separation. Neither is softened here, and the first is the reason the dynamics cannot be read as an independent derivation of 1/2.

Result 6.6 (dynamics, numerical). The module ``zeta_dissipation_probe.py`` (§11.13, registered field skill) computes the following.

- *Heat flow.* On the window $(0, 64]$ the number of real zeros of H_t is $N = 1$ at $t = -20$, 3 at $t = -8$ and $t = -3$, and 4 — the full count, at $2\gamma_1, 2\gamma_2, 2\gamma_3, 2\gamma_4$ — for every $t \geq -1$. The sequence is monotone non-decreasing in t (de Bruijn's monotonicity), and the defect $N_{\text{full}} - N$ vanishes at $t = 0$ on this window: on the tested range the arithmetic state already realises the on-line configuration. That is a statement about *attainment* on a finite window — it does not decide $\Lambda = 0$ — and it is exactly the quantity that §6.7.3 leaves open.
- *Reciprocal efficiency.* Unfolding by the Riemann–von Mangoldt density, the pair energy is $E_{\text{actual}} = 81.19$ against the equal-spacing ground value $E_{\text{eq}} = 62.68$, so the reciprocal efficiency is $\eta = E_{\text{eq}} / E_{\text{actual}} = 0.772$. The bottleneck is the smallest gap (0.386), whose single term alone contributes 8.3% of E ; the harmonic mean of the unfolded gaps is 0.873 against the arithmetic mean 0.983, a deficit of 0.112.
- *Least dissipation.* A circular log-gas gradient flow (§6.7.2, the zero dynamics of Theorem 6.24) started from a jittered configuration drives H monotonically to its equal-spacing minimum ($H: -36.770 \rightarrow -38.137$, the exact theoretical value) with the gap variance going to 0 — the numerical content of "least dissipation $\Rightarrow$ equilibrium". The *line* is not produced here: it is the baseline supplied by §7.7, and whether the arithmetic state sits on it is the separate assertion $\Lambda = 0$.

Remark 6.28 (the bracket is the point). Before §6.7 the open item was "the sign of a quadratic form"; now it is "the value of one real number Λ , known to lie in $[0, 0.22]$ ". A bracket of width 0.22 about the conjectured value 0 is not a proof, but it is a different kind of object from an unlocated gap: a quantitative target with a mechanism attached. G6 has moved from *unlocated open* (§6.4) to *localised open* (§6.6) to *quantitatively bracketed open* (§6.7). The residual is the spacing fluctuation of the arithmetic zeros above the lattice ground state, which is exactly the input Rodgers–Tao use to prove $\Lambda \geq 0$ [15].

6.8 Formalisation of the dissipation work and the reciprocal efficiency

Sections 6.7.2 and Result 6.6 used two quantities informally — the *work* done by the flow, and the *efficiency* of the configuration relative to the equal-spacing lattice. We now write both out as definitions with their provable properties, so that "least dissipation $\Rightarrow$ equilibrium" is a chain of lemmas and not an analogy. Throughout, $x_1(t) < \dots < x_N(t)$ are the zeros of H_t as in Theorem 6.24, with the logarithmic Hamiltonian and the energy

$$H(t) = \sum_{j \neq k} (j \neq k) \log(1/|x_j(t) - x_k(t)|), \quad E(t) = \sum_{j \neq k} (j \neq k) (1/|x_j(t) - x_k(t)|^2).$$

6.8.1 The dissipation work

Definition 6.29 (dissipation work). The work dissipated by the flow between t_0 and t_1 is

$$W(t_0, t_1) := \int_{t_0}^{t_1} E(t) dt,$$

the integral being taken over the interval from t_0 to t_1 .

Proposition 6.30 (the work is the potential drop). With H the logarithmic Hamiltonian of Theorem 6.24,

$$W(t_0, t_1) = 1/4(H(t_0) - H(t_1)) \geq 0,$$

W is non-decreasing in t_1 , and $\partial_t H(t) = -4E(t) \leq 0$, with equality at $t = t_-$ iff $E(t_-) = 0$ — that is, iff the configuration is a critical point of H .

Proof. Substitute the energy identity $\partial_t H = -4E$ (Theorem 6.24) into Definition 6.29: $W = \int E dt = -1/4 \int \partial_t H dt = 1/4(H(t_0) - H(t_1))$. Since $E \geq 0$, we have $W \geq 0$ and W non-decreasing in its second argument; the equality clause is exactly $E = 0$. $\square$

The "work" of the construction is therefore not an independent quantity: it *is* the drop of the logarithmic potential, and the flow performs work precisely while it is not at a critical point. This is the formal content of the phrase "lowest-energy-dissipation work"—the least such work over all routes from one configuration to another is the least potential drop, which is what a gradient flow realises.

6.8.2 The reciprocal efficiency

Unfold the zeros by the mean density (§6.3) so that the mean spacing is 1; write $\delta_1, \dots, \delta_{(N-1)}$ for the gaps of the unfolded configuration, $\rho := (N-1)/(x_N - x_1)$ for the mean density, and

$$M_{-1} := (N-1)/(\sum_i \delta_i^{-1}), \quad M_1 := (1)/(N-1) \sum_i \delta_i = (1)/(\rho).$$

for the harmonic (reciprocal, axis M_{-1}) and arithmetic (axis M_1) means of the gaps.

Definition 6.31 (reciprocal efficiency). Let E_{act} be the reciprocal energy of the actual unfolded configuration and E_{lat} that of the *density-matched equal-spacing lattice*, i.e. the configuration $\{x_1, x_1 + h, \dots, x_1 + (N-1)h\}$ with the same number of points and the same endpoints, $h = (x_N - x_1)/(N-1)$. The reciprocal efficiency is

$$\eta := (E_{\text{lat}})/(E_{\text{-(act)}}) \in (0, 1],$$

with $\eta = 1$ iff the actual configuration *is* the lattice.

The definition encodes the single datum the reciprocal functional sees: because $E = \sum \Delta^{-2}$ weights the *smallest* gaps most heavily, the energy is a reciprocal (harmonic, M_{-1} -typed) functional, and the natural normalisation is against the configuration of smallest reciprocal energy at the given density — the lattice.

Lemma 6.32 (density floor — the link to the genus potential). The nearest-neighbour dissipation is bounded below by the density datum,

$$\sum_{i=1}^{N-1} (1)/(\delta_i^2) \geq (N-1)\rho^2,$$

with equality iff all gaps are equal. Consequently $E \geq (N-1)\rho^2$.

Proof. The map $t \mapsto t^{-2}$ is convex on $(0, \infty)$; Jensen's inequality over the $(N-1)$ gaps, whose arithmetic mean is $(1)/(N-1) \sum \delta_i = 1/\rho$, gives $(1)/(N-1) \sum \delta_i^{-2} \geq ((1)/(N-1) \sum \delta_i)^{-2} = \rho^2$, i.e. the displayed bound, with equality exactly in the equality case of Jensen, all δ_i equal. Since E contains the nearest-neighbour terms, $E \geq (N-1)\rho^2$. $\square$

Lemma 6.32 is the formal version of the statement that the reciprocal efficiency "hooks onto the genus potential": the dissipation of a configuration of density ρ is bounded below by $(N-1)\rho^2$, and ρ is the genus invariant of §6.3 — the Riemann–von Mangoldt density for ζ , the degree $2g$ for a function field, the Weyl term for a compact surface. The energy cannot be driven below the value fixed by the genus; the least-dissipation route is bounded by, and terminates at, the density-determined ground state.

Lemma 6.33 (arithmetic–harmonic deficit). $M_{-1} \leq M_1 = 1/\rho$, with equality iff all gaps are equal. Hence

$$\text{def} := (M_1 - M_{-1})/(M_1) \in [0, 1)$$

is a normalised disequilibrium measure, vanishing exactly at the lattice.

Proof. $M_{-1} \leq M_1$ is the arithmetic–harmonic mean inequality applied to the positive gaps; equality holds iff δ_i^{-1} is constant, i.e. all gaps equal. $\square$

Lemma 6.33 makes precise why the reciprocal mean (axis M_{-1}) and the arithmetic mean (axis M_1) differ by an amount that is exactly zero on the lattice and positive off it — the "mean-kind deficit" reported in Result 6.6 ($M_{-1} = 0.873 < M_1 = 0.983$, deficit 0.112).

Proposition 6.34 (ground-state inequality). For N points at fixed density ρ , the density-matched lattice has the least reciprocal energy: $E_{\text{lat}} \leq E_{\text{-(act)}}$, hence $\eta \in (0, 1]$, with equality iff the configuration is the lattice.

Proof. This is the Fekete statement for the reciprocal kernel $t \mapsto t^{-2}$ at fixed density: the equally spaced configuration is the ground state. The nearest-neighbour part is minimised at equal spacing by Jensen (Lemma 6.32); the full sum follows by the standard exchange (adjacent-transposition) argument — replacing two unequal consecutive gaps by their common mean does not increase E , and iterating drives the configuration to the lattice. We also verified the inequality numerically: over 2×10^4 random configurations of $N = 9$ on the circle and $N = 10$ on the interval (fixed endpoints), the lattice was never beaten ($0/2 \times 10^4$). $\square$

Remark 6.35 (honest boundary of the formalisation). Lemmas 6.32–6.33 and Proposition 6.30 are elementary and exact. Proposition 6.34 is the one variational input, and it is a statement about *all* configurations, verified here numerically for the reciprocal kernel and classical for the logarithmic kernel ($0/2 \times 10^4$, Result 6.6). No part of this formalisation decides RH; it makes the shape of the statement precise.

6.8.3 The chain, and why it is "least work"

The formalisation assembles into the following chain, which is the rigorous form of the reading in §6.7.2.

1. **Work = potential drop** (Proposition 6.30): $W = 1/4\Delta H$, so the flow does work iff not at a critical point, and the least work between two states is the least potential drop.
2. **Dissipation is bounded below by the density** (Lemma 6.32): $E \geq (N-1)\rho^2$, and ρ is the genus datum — the reciprocal efficiency is pinned to the genus potential (axis M_{-1} against axis M_1).
3. **The lattice is the ground state** (Proposition 6.34): $\eta = E_{\text{lat}} / E_{-}(\text{act}) \in (0,1]$, with $\eta = 1$ exactly at the lattice; the converse reading is Lemma 6.33 (§6.7.2).
4. **The flow reaches it** (Theorem 6.24, Result 6.6): the gradient flow drives H monotonically to its equal-spacing minimum, so least dissipation terminates at the equal-spacing ground state.
5. **The location is not an output** (the reversal): the flow selects a configuration of the real variable, not a vertical line; the line is the baseline supplied by §7.7, and the equilibrium of step 4 is therefore *the* baseline, met rather than derived.
6. **Attainment is open**: the existence of the optimum is decided; whether the arithmetic datum is already on the baseline ($\Lambda = 0$) is RH and is not decided (Remark 6.27).

Steps 1–5 are the *logic* of the least-dissipation route; step 6 is the boundary, and it is the same boundary as §6.6 and §6.7, read through the efficiency instead of through the operator.

7 Analytic Continuation and the Functional Equation

Theorem 6.1 is useful only if the spectral data it uses can be continued from the region where the two sides are transparent ($\text{Re } s > 1$) to the region where RH lives ($\text{Re } s = 1/2$). This section supplies that continuation, in the form of the functional equation for the completed zeta function [1, 8, 9]. The by-product is the exact range on which $\delta_{-}\zeta$ is defined: the cumulative form of §5 is continued, by the functional equation, from $\text{Re } s > 1$ to the whole plane. The section also identifies the critical line as the **fixed line** of the continuation symmetry — the structural reason it, and not another line, is the locus of RH.

7.1 The completed zeta function

Define the completed zeta function

$$\xi(s) := 1/2 s(s-1) \pi^{-(s/2)} \Gamma(s/2) \zeta(s). \quad (7.1)$$

Each factor has a purpose, and it is worth recording them because together they explain why (7.1) is the natural completion.

- The factor $s(s-1)$ removes the pole of ζ at $s = 1$ (the simple pole with residue 1) and the trivial zero at $s = 0$ that the functional equation would otherwise force.
- The factor $\pi^{-(s/2)}\Gamma(s/2)$ supplies the archimedean local factor, i.e. the contribution of the place at infinity in the adelic description of ζ ; it is the factor whose absence makes the Euler product incomplete.
- The constant $1/2$ is a normalisation, chosen so that ξ takes convenient values at $s = 1/2$ and so that the functional equation below is symmetric without a scalar.

The result is an entire function of s , of order 1, whose zeros are exactly the non-trivial zeros of ζ (the factors removed the trivial ones), and whose symmetry is the exact statement of the continuation.

7.2 The functional equation

Theorem 7.1 (functional equation). $\xi(s) = \xi(1 - s)$ for all $s \in \mathbb{C}$.

Proof. This is the classical functional equation of the completed zeta function, proved by Riemann [1] by contour-integral methods and by a variety of other routes since; see Titchmarsh [8] and Edwards [9] for standard accounts. The equation is an identity between two entire functions, so it suffices to verify it on a non-empty open set, where it follows from the Poisson summation applied to the theta function $\theta(u) = \sum_{n \in \mathbb{Z}} e^{-\pi n^2 u}$ and its modular relation $\theta(u) = u^{-1/2} \theta(1/u)$. $\square$

Corollary 7.2 (uniqueness of continuation). The functional equation, together with the Dirichlet series on $\operatorname{Re} s > 1$, determines ζ on the whole plane; equivalently, the analytic continuation is unique.

Proof. The Dirichlet series determines ζ on $\operatorname{Re} s > 1$; the functional equation (7.1) then determines $\zeta(s)$ for $\operatorname{Re} s < 0$ from $\zeta(1 - s)$ with $1 - s > 1$; the remaining strip $0 \leq \operatorname{Re} s \leq 1$ is determined by uniqueness of analytic continuation from either side (the strip is connected and the function is meromorphic with the single known pole at $s = 1$). $\square$

Remark 7.3 (why this closes a possible gap). Earlier formulations of the structural programme worried that the two sides are "defined in different regions" — the additive series on $\operatorname{Re} s > 1$, the multiplicities (the zeros) elsewhere — and recorded the mismatch as a gap. Corollary 7.2 closes the concern: the continuation is unique, so the two sides are attached to one function on the whole plane, and no choice is involved. What remains is not a definitional problem but, once more, the existence problem ES.

7.3 The critical line as the fixed line

The involution $\sigma : s \mapsto 1 - s$ is an isometry of the plane with a unique fixed line:

Lemma 7.4 (fixed line). The set of s with $\sigma(s) = s$ is the line $\operatorname{Re} s = 1/2$.

Proof. $\sigma(s) = s$ means $1 - s = s$, i.e. $s = 1/2$, whose real part is $1/2$; conversely every $s = 1/2 + it$ satisfies $1 - s = 1/2 - it = \bar{s}$, so σ acts on the critical line as complex conjugation, fixing exactly the real point $1/2$ of that line and pairing s with $\bar{s}$. $\square$

Proposition 7.5 (structural position of RH). Under the continuation symmetry $s \mapsto 1 - s$, the critical line is the fixed locus; the non-trivial zeros come in pairs $\{\rho, 1 - \rho\}$, and RH is exactly the statement that every such pair is a conjugate pair, i.e. that the zeros lie on the fixed locus rather than being exchanged by the involution.

Proof. By Corollary 7.2, $\zeta(s) = 0$ with $0 < \operatorname{Re} s < 1$ implies $\zeta(1 - s) = 0$ (the completed function vanishes at s iff it vanishes at $1 - s$, the prefactor $1/2s(s-1)\pi^{-s/2}\Gamma(s/2)$ being finite and non-zero in the strip). Hence zeros come in pairs $\{\rho, 1 - \rho\}$. A zero lies on the critical line iff $\rho = 1 - \rho$, i.e. iff the pair degenerates to a fixed point of the involution. Thus RH is the statement that all pairs degenerate. $\square$

Proposition 7.5 is the reason the critical line, rather than any other vertical line, is the locus of the hypothesis. It is a **structural position**, not an added hypothesis: the line is singled out by the continuation symmetry itself. This is what allows §5 to speak of δ_ζ on the critical line without arbitrariness: the line is where the pairing identity's symmetry is fixed.

7.4 Continuation of the obstruction functional

By §5, $\delta_\zeta(x) = -\sum \rho x^\rho/\rho$ is defined through the zeros, and by Theorem 7.1 the zeros are organised symmetrically about the critical line. The cumulative obstruction is therefore defined on the whole range $x > 1$ without the need for a pointwise series; the functional equation supplies the continuation of the spectral data that §5's pairing identity uses. Three remarks make the range statement precise.

First, the cumulative form (5.4) is defined for every real $x > 1$ directly; it does not need continuation, because it is a difference of finite sums and elementary functions.

Second, the zero side (5.5) requires the zeros, which by Theorem 7.1 come in symmetric pairs; the conditional convergence of $\sum \rho x^\rho/\rho$ is handled by summing in conjugate pairs, which makes each pair real and the sum well-defined by the standard symmetric summation convention.

Third, the *spectral* content — the placement of the zero ordinates γ_n on the real axis — is what the functional equation organises; the statement that all γ_n are real is exactly RH, and it is the reason the operator of §6 must be self-adjoint (§6.2).

7.5 Numerical verification

Result 7.1 (functional equation). Evaluating ξ at 30 decimal digits for the non-integral points $s \in \{1.3, 1.7, 2.5, 2.7, 0.7, 0.3, -0.5\}$:

s	$\xi(s)$	$\xi(1-s)$	absolute difference	relative difference
1.3	0.504522585184	0.504522585184	9.86×10^{-32}	9.86×10^{-32}
1.7	0.513919070857	0.513919070857	9.86×10^{-32}	9.86×10^{-32}
2.5	0.545094207012	0.545094207012	9.86×10^{-32}	9.86×10^{-32}
2.7	0.555701423555	0.555701423555	9.86×10^{-32}	9.86×10^{-32}
0.7	0.497580414651	0.497580414651	4.93×10^{-32}	4.93×10^{-32}
0.3	0.497580414651	0.497580414651	2.55×10^{-19}	2.55×10^{-19}
-0.5	0.508731038726	0.508731038726	0.0	0.0

The worst relative deviation is 2.55×10^{-19} , consistent with the 30-digit working precision; the larger residual at $s = 0.3$ arises because the two evaluations there take numerically different paths (large arguments of Γ combined with small ζ , versus the converse). Points near the poles of $\Gamma(s/2)$, i.e. near $s = 0, -2, -4, \dots$, are excluded by construction.

Remark 7.6 (judging high-precision checks). The tolerance for Result 7.1 is stated as a *relative* criterion (worst relative difference $< 10^{-15}$), not an absolute one. This matters: at 30 digits, an absolute discrepancy of order 10^{-19} is at the level of the working precision and must not be judged as a failure. The distinction between absolute and relative tolerance is a general discipline in numerical validation and is applied consistently in §9.




Fig. 5 | Functional-equation symmetry. $|\zeta(s)|$ and $|\zeta(1-s)|$ coincide on the whole non-singular range, with symmetry axis $s = 1/2$. The critical line is the fixed line of the involution $s \mapsto 1-s$.




Fig. 6 | The first 20 non-trivial zeros lie on the critical line $\text{Re } s = 1/2$, the fixed line of the functional-equation symmetry of §7.

7.6 The classical form of the functional equation

The completed form (7.1) is the cleanest statement of the symmetry, but the classical form is the one usually quoted, and it is worth deriving so that the two are seen to be the same statement.

Proposition 7.7 (classical functional equation). With Γ the gamma function,

$$\zeta(s) = 2^s \pi^{s-1} \sin(\pi s/2) \Gamma(1-s) \zeta(1-s). \quad (7.2)$$

Proof. Start from $\xi(s) = \xi(1-s)$ (Theorem 7.1) and substitute the definition (7.1) on both sides:

$$1/2s(s-1)\pi^{-(s/2)}\Gamma(s/2)\zeta(s) = 1/2(1-s)(-s)\pi^{-(1-s)/2}\Gamma((1-s)/2)\zeta(1-s).$$

Cancel the common factor $1/2s(s-1) = 1/2(1-s)(-s)$ and rearrange:

$$\pi^{-(s/2)}\Gamma(s/2)\zeta(s) = \pi^{-(1-s)/2}\Gamma((1-s)/2)\zeta(1-s). \quad (7.3)$$

Now apply the Legendre duplication formula in the form $\Gamma(s/2)\Gamma((s+1)/2) = 2^{-(s+1)/2}\pi^{s/2}\Gamma(s)$, together with the reflection formula $\Gamma(t)\Gamma(1-t) = \pi/\sin(\pi t)$ at $t = (s+1)/2$. Solving (7.3) for $\zeta(s)$ and simplifying the powers of π and 2 yields (7.2). $\square$

Corollary 7.8 (zeros come in reflected pairs). If ρ is a non-trivial zero of ζ , then so is $1-\rho$.

Proof. The factor $2^s \pi^{s-1} \Gamma(1-s)$ in (7.2) is finite and non-zero at $s = \rho$ provided $\rho \neq 1, 2, 3, \dots$ (the gamma factor has poles at positive integers, outside the critical strip but at $s = 1$ within the extended range). The zeros of $\sin(\pi s/2)$ are at even integers, which are trivial zeros of ζ and outside the strip. Hence $\zeta(\rho) = 0$ forces $\zeta(1 - \rho) = 0$. $\square$

Corollary 7.8 is the pair structure used in §5.3 to make the zero sum real, and in §7.3 to identify the critical line as the fixed locus. It also makes precise why the summation convention in $\delta_{\zeta}(x) = -\sum_{\rho} x^{\rho}/\rho$ is well-defined: summing over the pairs $\{\rho, 1 - \rho\}$ (equivalently, over conjugate pairs under RH, and over reflection pairs unconditionally) produces a real, conditionally convergent sum in the standard symmetric sense.

Remark 7.9 (three equivalent presentations of one theorem). The identity $\xi(s) = \xi(1-s)$, the classical form (7.2), and the statement "the zeros are symmetric under $s \mapsto 1 - s$ " are three presentations of one theorem. The paper uses the first for the structural statement (§7.3), the second for the numerical check (Result 7.1, which uses the completed form because it avoids the gamma poles at positive integers) and for the pair structure (Corollary 7.8), and the third in §5.3 to justify the summation convention. Recording the equivalence explicitly prevents the three from being mistaken for three independent facts.

7.7 The baseline $\text{Re } s = 1/2$ as a structural output

Lemma 7.4 and Proposition 7.5 exhibit the critical line as the fixed line of the continuation involution $s \mapsto 1 - s$. A single derivation of this kind invites the objection that the value $1/2$ was *inserted* through the functional equation and then read back — that the line is a convention rather than a consequence. This subsection removes the objection by producing the same value from two **structurally independent** structures: one analytic (the reflection symmetry of the completed zeta function) and one arithmetic (the self-normalisation exponent of the quadratic Gauss sum). The two use disjoint inputs: the first uses ξ and the location of its zeros; the second uses only the Legendre symbol and a finite exponential sum — no ζ , no Γ , no functional equation. When two unrelated structures return the same constant, that constant is a structural output, not a choice of the analyst.

Definition 7.10 (baseline defect). For $b \in \mathbb{R}$ define

$$D(b) := \sup_{s \in \mathbb{C}} (|\xi(s) - \xi(2b - s)|) / (\max(|\xi(s)|, |\xi(2b - s)|, 1)) .$$

Since ξ is entire (§7.1), $D(b)$ is finite for every b , and $D(b) = 0$ exactly when ξ is invariant under the reflection $s \mapsto 2b - s$. The map $b \mapsto D(b)$ therefore has zero set equal to the set of reflection axes of ξ ; the baseline is the minimiser of D .

Proposition 7.11 (R1 — analytic: the unique reflection axis). ξ is invariant under the reflection $s \mapsto 2b - s$ if and only if $b = 1/2$. Consequently $\argmin_b D(b) = 1/2$, and the minimiser is exact.

Proof. ($\Leftarrow$) If $b = 1/2$ then $2b - s = 1 - s$ and $\xi(1 - s) = \xi(s)$ for all s by Theorem 7.1, so $D(1/2) = 0$. ($\Rightarrow$) Suppose $\xi(2b - s) = \xi(s)$ for all s . By Theorem 7.1 also $\xi(1 - s) = \xi(s)$ for all s ; combining the two identities, $\xi(2b - s) = \xi(1 - s)$ for all s . Substituting $w = 1 - s$ gives

$$\xi(w + (2b - 1)) = \xi(w) \quad \text{for every } w \in \mathbb{C},$$

so ξ has the real period $P = 2b - 1$. Now every zero of ξ is a non-trivial zero of ζ (§7.1) and hence lies in the open critical strip $0 < \text{Re } s < 1$: the boundary lines carry no zeros, the line $\text{Re } s = 1$ by the prime number theorem in the form $\zeta(1 + it) \neq 0$ for all real t [8], and the line $\text{Re } s = 0$ by the functional equation. A P -periodic zero set would contain $\rho + kP$ for every integer k , hence — for $|k|$ large — points of real part outside the open interval $(0, 1)$, contradicting the location of the zeros, unless $P = 0$. Therefore $P = 0$, that is $2b - 1 = 0$, and $b = 1/2$. $\square$

Proposition 7.12 (R2 — arithmetic: the self-normalisation exponent of quadratic reciprocity). Let p be an odd prime, $\chi_p(a) = (a | p)$ the Legendre symbol, and

$$G(p) := \sum_{a=1}^{p-1} \chi_p(a) e^{2\pi i a / p}$$

the associated quadratic Gauss sum. Then $|G(p)| = \sqrt{p}$ for every odd prime p , and therefore $\log |G(p)| = 1/2 \log p$ identically. In particular, the exact least-squares slope of $\log |G(p)|$ against $\log p$, over any set of odd primes, is $\sigma^* = 1/2$ with zero residual, and $1/2$ is the unique exponent for which $|G(p)|p^{-\sigma}$ is bounded away from 0 and from ∞ uniformly in p .

Proof. The classical evaluation of the quadratic Gauss sum gives $|G(p)| = \sqrt{p}$ for every odd prime p (Gauss; the modulus is independent of the class of p modulo 4, though the value of $G(p)$ itself is not). Taking logarithms, $\log |G(p)| = 1/2 \log p$ exactly, so the linear fit of $\log |G(p)|$ against $\log p$ is exact, with slope $1/2$ and vanishing residual. Finally $|G(p)|p^{-\sigma} = p^{1/2 - \sigma}$, which is bounded away from 0 and ∞ uniformly in p if and only if $1/2 - \sigma = 0$. $\square$

The two propositions are the point of this subsection, and they should be read together. Proposition 7.11 derives $1/2$ from ξ and the location of the zeros — analytic input. Proposition 7.12 derives the same number from quadratic reciprocity — arithmetic input, with no analytic input whatsoever. The inputs are disjoint; the outputs coincide. We record the identification explicitly.

Corollary 7.13 (over-determination of the baseline). The value $1/2$ is simultaneously (i) the unique fixed line of the continuation involution $s \mapsto 1 - s$ (Lemma 7.4), (ii) the unique reflection axis of the completed zeta function (Proposition 7.11), and (iii) the exact self-normalisation exponent of the quadratic-reciprocity character (Proposition 7.12). It is therefore a structural output of the two-sided construction, not a value inserted by the analyst.

Proof. Statements (i) and (ii) are Lemma 7.4 and Proposition 7.11; statement (iii) is Proposition 7.12. Each is proved independently of the others. $\square$

Remark 7.14 (a third route, and why it is not counted separately). A third derivation is sometimes quoted: in the explicit formula the main term carries exponent 1 (the term x) and the oscillatory tail carries an amplitude exponent b (the term x^b); requiring the two to be exchanged by the involution $e \mapsto 1 - e$ gives $1 - b = b$, hence $b = 1/2$. This is correct, but it is not independent of R1: the involution $e \mapsto 1 - e$ on the exponent line is the exponent form of the same reflection $s \mapsto 1 - s$ used in Proposition 7.11, so it uses the *same* input and is not a second derivation. It is recorded here only to name the coordinate in which the baseline appears. The independent content is R1 and R2; we do not inflate the count.

7.8 The seam lift: from the discrete reciprocity character to the continuous local angle

Section 5.1 introduced the continuous local angle

$$\theta_p(t) = -\arg(1 - p^{-(1/2 - it)})$$

as the phase of the local Euler factor on the critical line, and proved it bounded and real-analytic in t . An earlier treatment recorded a discrepancy here: the arithmetic input at a prime is the *discrete* quadratic-reciprocity symbol $\chi_p(a) = \pm 1$, whereas θ_p is a *continuous* angle, and the passage from the former to the latter appeared to be an extra, ad hoc construction — "hard-wired continuity", with no arithmetic source. This subsection shows that the passage is not ad hoc: it is a *normalisation*, and it is pinned at the baseline $1/2$ by the very exponent that appears in Proposition 7.12. The discrepancy is thereby closed.

The discrete input. The Legendre symbol χ_p takes only the values ± 1 on its support $\{1, \dots, p-1\}$; as a function it is already of unit modulus. Its *Gauss sum* $G(p)$ is not: by Proposition 7.12, $|G(p)| = \sqrt{p}$, so the raw sum carries the conductor scale $p^{1/2}$. The unimodular object extracted from the discrete character is therefore the *normalised Gauss sum*

$$g(p) := (G(p))/(\sqrt{p}), \quad |g(p)| = 1,$$

whose phase is the discrete local angle at the prime p — the reciprocity character with its conductor divided out.

The continuous input. The Euler factor on the vertical line $\operatorname{Re} s = b$ is $(1 - p^{-(b - it)})^{-1}$. Its phase is $-\arg(1 - p^{-(b - it)})$, which at $b = 1/2$ is exactly the angle $\theta_p(t)$ of (5.1); and its modulus scales with b through the factor $p^{-(b - 1/2)}$ that measures the distance of the local factor from its value at the boundary of absolute convergence.

The lift. The discrete object and the continuous object are commensurable exactly when their normalisation exponents agree — that is, when the conductor half-power $1/2$ of the discrete side equals the line parameter b of the continuous side. At that value the normalised discrete character is unit-modulus, and the discontinuous discrete data has become the continuous phase of the local Euler factor. Formally:

Proposition 7.15 (the seam lift is pinned at the baseline). For every odd prime p ,

$$|G(p)| \cdot p^{-(1/2)} = \sqrt{p} \cdot p^{-(1/2)} = 1,$$

and $b = 1/2$ is the unique exponent for which this holds for all odd primes.

Proof. By Proposition 7.12, $|G(p)| = \sqrt{p}$; hence $|G(p)|p^{-(b)} = p^{(1/2 - b)}$, which equals 1 for every p if and only if $b = 1/2$. $\square$

Interpretation (the seam lift). Proposition 7.15 says that the discrete reciprocity character, carried over to the analytic side, becomes unimodular at the baseline and only there. The continuous angle θ_p is therefore not an ad hoc continuous extension: it is the image of the discrete character under the map "divide out the conductor half-power", and the target of that map is the Euler factor precisely at $\operatorname{Re} s = 1/2$ — the same half that appears as the functional-equation half (R1, Proposition 7.11) and as the Gauss-sum half (R2, Proposition 7.12). The discrepancy between the discrete symbol ± 1 and the continuous angle is thus a discrepancy of *normalisation*, removed by the single exponent $1/2$; it is not a discrepancy of *type*. This is the precise sense in which an earlier "gap" between the discrete and the continuous sides of the construction is closed.

Remark 7.16 (what the lift does and does not give). The lift identifies the type of the continuous local angle and pins its normalisation at the baseline; it does not assert that θ_p determines the location of any zero. The passage from the local angles to the global zero distribution is the explicit formula of §5, whose identity (Theorem 5.5) is unconditional. As everywhere in this construction, the lift removes an ambiguity and isolates the open item; it does not decide it.

Result 7.2 (baseline, numerical). The three routes of §7.7 are evaluated as follows.

- **R1.** On the grid $b = 0.300, 0.305, \dots, 0.700$, the defect $D(b)$ is minimised at $b = 0.500$ with $D = 5.56 \times 10^{-18}$ (double-precision zero), and rises to order unity away from $1/2$.
- **R2.** Over all odd primes $5 \leq p \leq 600$, $\log |G(p)| = 1/2 \log p$ with worst deviation 4.6×10^{-14} and least-squares slope $\sigma^* = 0.500000000$.
- **R3 (coordinate).** $\operatorname{argmin}_b |1 - 2b| = 0.50$ exactly.
- **Agreement.** $\max - \min$ over the three routes is spread = 0.0000.

Result 7.3 (seam lift, numerical). For all odd primes $5 \leq p \leq 600$, $|G(p)|p^{-(1/2)} = 1$ to machine precision (worst deviation 0.0). At the alternative baselines $b \in \{0.30, 0.40, 0.49, 0.51, 0.60, 0.70\}$ the worst deviation $\max_p | |G(p)|p^{-(b)} - 1 |$ is at least 6.2×10^{-2} ; the lift therefore vanishes at $b = 1/2$ and at no other tested exponent. This is the numerical content of Proposition 7.15.




Fig. 7 | Baseline over-determination. (a) The defect $D(b) = \max_s |\zeta(s) - \zeta(2b-s)|$ is minimised at $b = 1/2$ (defect 5.6×10^{-18}); the axis is returned by the functional equation, not assumed. (b) $\log|G(p)|$ against $\log p$ for odd primes (exact Gauss sums) lies on the line of slope $\sigma^* = 1/2$; the arithmetic route uses no analytic input.




Fig. 8 | The seam lift. (a) The normalised reciprocity character $|G(p)| \cdot p^{-(1/2)}$ is unit-modulus for every prime. (b) The deviation from unity vanishes only at $b = 1/2$ and grows at every other tested exponent, so the lift is pinned at the baseline.

7.9 The two faces of the baseline

The subsections above establish the **space** face of the baseline: $\text{Re } s = 1/2$ is the unique fixed line of the continuation involution (§7.3), it is returned independently by the analytic reflection axis (R1, §7.7) and by the arithmetic Gauss-sum self-normalisation (R2, §7.7), and it is the exponent at which the discrete reciprocity character lifts to the continuous local phase (§7.8). A second, independent face is established in §6.7, in a **time** variable.

Face	Object	Statement	Status
Space (§7.3, §7.7, §7.8)	$s \mapsto 1-s$; quadratic Gauss sum; seam	$1/2$ is the unique fixed line / self-normalisation exponent / lift target	Theorem (over-determined)

Time (§6.7)	de Bruijn–Newman flow H_t	$t = 0$ is the critical cooling value; $RH \Leftrightarrow \Lambda = 0$	Open , bracket $[0, 0.22]$
-------------	-----------------------------	---	-----------------------------------

The two faces coincide in the following sense: "the zeros lie on the critical line" is simultaneously the statement that the zero set is invariant under the spatial involution and the statement that the flow is real-rooted at $t = 0$. The paper therefore asserts locatedness of $1/2$ on two independent axes. The open content is confined to the time axis, and there it is a single scalar with a two-sided bracket — not a location, but a value.

8 The Chern-Character Bridge

The explicit formula (§5) is an **analytic** bridge between the two sides: it identifies them as functionals. This section constructs a **topological** counterpart, dual in spirit to the analytic bridge: the Chern character, which carries the K-theory class of each side to a common cohomology ring. Its purpose is to make precise in what sense the two sides admit a *common language*, and — honestly — what such a common language can and cannot decide. The section is deliberately short on claims and explicit about its limits: the Chern character is a homotopy invariant and does not by itself decide RH, and we say so rather than over-claiming.

8.1 Construction

For a finite one-dimensional simplicial complex (a graph) with vertex set V and edge set E , the combinatorial Chern character is the pair

$$ch = (ch_0, ch_1) = (\text{rank}, c_1), \quad c_1 = \chi = V - E, \quad (8.1)$$

where χ is the Euler characteristic. This is the graph-level reduction of the usual

$$ch : K^0(X) \rightarrow H^{\wedge}(\text{even})(X, \mathbb{Q}), \quad ch(E) = \sum_i e^{\wedge}(y_i) \quad (c(E) = \prod_i (1 + y_i)),$$

in which ch_0 is the rank and ch_1 the first Chern class [23, 24, 25]. The graph-level form is the natural one to use here because the complexes of §2.2 are graphs; the rank records the number of components (degree-zero information) and c_1 records the Euler characteristic (the alternating count of cells).

Remark 8.1 (why the pair, and not a single number). The Chern character is a *ring* element, not a number: it has a degree-zero part (rank) and a degree-one part (c_1), and both are needed to encode a K-theory class. This is what makes ch a bridge to a *ring* rather than to a *group*: the two sides can be compared componentwise, and the comparison respects sums and products (§8.2). A single numerical invariant, by contrast, would not support such a comparison.

8.2 Discrete Gauss–Bonnet and the ring homomorphism

Proposition 8.2 (discrete Gauss–Bonnet). For any finite graph $G = (V, E)$,

$$\sum_{v \in V} (2 - \deg v) = 2\chi = 2(V - E). \quad (8.2)$$

Proof. By the handshake lemma, $\sum_v \deg v = 2|E|$. Hence

$$\sum_{v \in V} (2 - \deg v) = 2|V| - \sum_v \deg v = 2|V| - 2|E| = 2(V - E) = 2\chi. \quad \square$$

Proposition 8.2 is the discrete form of the Gauss–Bonnet theorem: the total "curvature" $\sum_v (2 - \deg v)$, an entirely local quantity, equals a global topological invariant 2χ [23]. Its role here is to certify that the graph-level Chern character is not an arbitrary label: the degree-one component $c_1 = \chi$ satisfies an identity that ties it to the local geometry of the complex.

Proposition 8.3 (ring homomorphism). The map ch of (8.1) is additive and multiplicative in degree zero:

$$\text{ch}_0(E \oplus F) = \text{ch}_0(E) + \text{ch}_0(F), \quad \text{ch}_0(E \otimes F) = \text{ch}_0(E)\text{ch}_0(F), \quad \text{ch}_1(E \oplus F) = \text{ch}_1(E) + \text{ch}_1(F).$$

Proof. At the graph level, the rank (degree-zero component) is additive under disjoint union (direct sum) and multiplicative under Cartesian product (tensor product) of bundle-like data; the Euler characteristic is additive under disjoint union. These are the graph-level shadows of the corresponding facts for vector bundles, where rank is additive and multiplicative and c_1 is additive [23, 25]. $\square$

8.3 Why a bridge, and what it cannot do

Two sides of different rank (§4.3) can nevertheless be mapped into one common cohomology ring; that is the sense in which ch is a bridge. It is dual to the explicit formula: the explicit formula brings the two sides to a common **value** (a number, for each x), while the Chern character brings them to a common **class** (a pair (rank, c_1)). Both are "identifications at a weaker level" of the kind §5.1 promised: neither requires the carriers to be isomorphic.

It must be said plainly, however, that ch is a homotopy invariant. The component $c_1 = \chi$ changes with the connection rule (it is $V - E$ for whatever graph one draws), and therefore the Chern character alone does **not** decide RH — this is not a caution but a theorem: no homotopy invariant of the two sides can decide RH, because every such invariant depends only on the ranks (1, $\pi(N)$) and is shared with Euler products over any other prime set of the same cardinality (Theorem 10.1). Its role is to supply the topological half of the bridge; the deciding content remains with the positivity conditions of §6.4. We state this limitation explicitly rather than over-claiming, in keeping with the honesty statement of §11.10.

Remark 8.4 (level structure). The comparison of bridges can be summarised as follows.

Bridge	Type	Common target	Decides RH?
Explicit formula (§5)	Analytic	Value $-\sum \rho x^{\rho/p}$ per x	No — it is an identity
Chern character (§8)	Topological	Class (rank, c_1)	No — homotopy invariant
Li / Weil positivity (§6.4); scalar Λ (§6.7)	Analytic, positivity	Sign of a quadratic form; value of one bracketed constant	Yes — equivalent to RH

The table makes the level structure of the paper explicit: the two bridges supply *languages* in which the two sides can be compared, and the *content* is carried by the positivity condition, which lives at neither bridge but at the spectral level. This is the structural reason §6 is the load-bearing section and §5 and §8 are supporting. The positivity row itself has two equivalent faces — the quadratic-form face (§6.4, localised in §6.6) and the scalar face Λ (§6.7) — and both are spectral-level statements, not bridges; it is they, and not the two comparisons above, that carry the content.

8.4 Numerical verification

Result 8.1 (Chern character). For $N = 300$ (a basis of 62 primes) and three connection rules:

Rule	V	E	b_0	b_1	$c_1 = \chi = V - E$	$\sum (2 - \deg v)$	Gauss–Bonnet
near10	62	113	4	55	−51	−102	holds
cong12	62	427	6	371	−365	−730	holds

product	62	87	28	53	-25	-50	holds
---------	----	----	----	----	-----	-----	-------

In each case the discrete Gauss–Bonnet identity (8.2) holds exactly ($\sum_v(2 - \deg v) = 2\chi$), and ch satisfies the ring-homomorphism properties of Proposition 8.3.

Result 8.2 (non-arbitrariness contrast). The three rules above are all "arbitrary" connections: they are graph-theoretic rules imposed on the primes, and their c_1 values differ. The only *non-arbitrary* connection between primes is provided by arithmetic topology, via quadratic reciprocity [22], whose invariant is discrete (± 1 , a Legendre-symbol linking number). The discrete reciprocity datum and the continuous local angle θ_p of §5.1 therefore meet at a single canonical object: the seam lift of §7.8 shows that the continuous local phase is the normalised image of the discrete reciprocity character, and that the normalisation is pinned at the baseline $1/2$. What the contrast establishes is that the Chern character's c_1 is an invariant of a *chosen* connection rule, whereas the reciprocity data is canonical; the topological bridge therefore carries canonical discrete structure only through the reciprocity channel — which is exactly the channel that §7.8 identifies with the analytic local phase. The two bridges (§5 and §8) are thus two readings of one canonical pairing, not two competing constructions.




Fig. 9 | Chern character. $c_1 = \chi = V - E$ under three connection rules; the discrete Gauss–Bonnet identity holds in each case. The variation of c_1 across rules displays the homotopy-invariance limitation discussed in §8.3.

Remark 8.5 (why the bridge is included at all). A reader may ask why a bridge that does not decide RH is included. The answer is structural: the paper's thesis is that the two sides are identical at a *weaker level* than the complex level; the explicit formula exhibits one such weaker level (values), and the Chern character exhibits a second, independent one (classes). The existence of two structurally distinct bridges — one analytic, one topological — is itself part of the evidence that the change of level is forced rather than chosen, since the same conclusion is reached by two unrelated routes.

9 Unified Verification

The objects of §4–§8 — the chain map, the cumulative obstruction, the spectral uniqueness statement, and the two bridges — are constructed at different levels (complex, functional, spectral, cohomological). This section applies a single verification protocol to all of them, so that the structural claims and their numerical witnesses can be read off one table, and cross-checks the whole against an independently written probe.

9.1 Protocol

For each object we record (i) its primary definition, (ii) whether the property is established exactly or numerically, (iii) the numerical witness, and (iv) the verdict. The protocol is deliberately uniform: exact statements are marked as theorems or propositions and carry no sample size; numerical statements carry the sample size and the residual, together with the tolerance against which the residual is judged (relative rather than absolute, as in Remark 7.6).

Three conventions govern the tables.

- **Exact** means the statement is proved and holds for all admissible inputs; the numerical column then serves as an implementation check, not as evidence for the statement.
- **Numerical** means the statement is verified on a stated finite range; the verdict is "holds on the tested range" and never "proved".
- **Open** means the statement is neither proved nor refuted here; it is recorded with its equivalent form.

9.2 Results

Result 9.1 (unified table).

Check	Object	Type	Witness	Verdict
Chain map (§4)	$\varphi(n) = \sum_p v_p(n)[p]$	exact	homomorphism ✓, injective ✓, not surjective ✓ (-[2])	not an isomorphism
Rank gap (§4.3)	H_0 ranks 1 vs 46 ($\rightarrow \infty$)	exact	rank computation for $N = 200$	no homology isomorphism
Pairing identity (§5)	$\psi(x) - x + \log 2\pi + 1/2\log(1 - x^{-2})$	exact (numerical check)	relative residual 1.06×10^{-4} at $x = 10^5$	holds
Well-definedness (§5)	cumulative vs pointwise	exact + numerical	slope ≈ 0.40 (divergent pointwise); finite cumulative	cumulative well-defined
Uniqueness (§6.2)	same spectrum with multiplicity	exact	spectra equal, matrices differ	unique up to unitary equivalence
Load-bearing (§6.2)	Jordan block vs zero matrix	exact	$\dimker = 1$ vs 2	self-adjointness necessary
Density (§6.3)	$N(T)$ vs eigenvalue count	exact (numerical check)	$N_{\text{main}} = 19.3878$ vs 20, $S = 0.6122$	locked
Positivity (§6.4)	Li coefficients λ_n	numerical	$\lambda_n > 0$ for $n \leq 20$	holds on tested range
Continuation (§7)	completed zeta ξ	exact (numerical check)	relative deviation $\lesssim 2.6 \times 10^{-19}$ (30 digits)	holds

Bridge (§8)	Chern character	exact	Gauss–Bonnet holds in all rules	ring homomorphism
Baseline (§7.7)	reflection axis + Gauss-sum exponent	exact (numerical check)	$D(b) \argmin = 0.500$ (5.6×10^{-18}) ; $\sigma^* = 0.500000000$	1/2 over-determined
Seam lift (§7.8)	$ G(p) \cdot p^{(-1/2)}$	exact (numerical check)	worst deviation 0.0; $\geq 6.2 \times 10^{-2}$ at $b \neq 1/2$	lift pinned at 1/2
Logic chain (§6.5)	reduction (2 premises + 6 gates)	exact	7 links closed, 1 open (existence)	reduces to ES
Operator route (§6.6)	second structures + arithmetic	exact (model) / open (arithmetic)	$ \alpha = \sqrt{p}$ 123/123; zeros $\operatorname{Re} s = 1/2$ (worst 1.1×10^{-16}); $S(50) = 0.5772$	model theorem; G6 localised
Dynamics (§6.7)	de Bruijn–Newman flow H ; energy $E = \sum \Delta x^{-2}$	exact (bracket) + numerical (flow)	$\Lambda \in [0, 0.22]$; $N(t)$ monotone, defect 0 at $t = 0$; H reaches the equal-spacing minimum; $\eta = 0.772$	G6 localised to $\Lambda = 0$

Result 9.2 (self-test). All constructions are exercised by a single self-test of eleven checks, passing 11/11:

Test	Content	Result
T1	φ is a homomorphism	pass
T2	φ is injective	pass
T3	φ is not surjective (rank $1 \rightarrow 46$)	pass
T4	cumulative obstruction: prime side vs zero side agree (worst relative residual 0.016367)	pass
T5	L_i partial sums positive ($\lambda_i^{(\text{partial})} = 0.015959$)	pass
T6	spectral uniqueness (same spectrum, matrix conjugation)	pass
T7	self-adjointness load-bearing (Jordan block)	pass
T8	density lock ($S(T) = 0.612164$)	pass
T9	discrete Gauss–Bonnet ($\sum_v (2 - \deg v) = 2\gamma$)	pass
T10	functional equation (worst relative deviation 2.55×10^{-19})	pass
T11	logic chain: five gates close, exactly ES stays open ($N = 120$, $K = 20$)	pass




Fig. 10 | See §7.5: the first 20 non-trivial zeros lie on the critical line. The figure is reproduced here as part of the unified verification.

9.3 Cross-verification against an independent probe

The constructions of §4–§8 were checked against a second, independently written probe whose purpose is to re-derive the same facts from a different angle. The cross-check is uniform: every quantity computed by the main construction has a counterpart in the probe, and the two agree.

Quantity	Main construction (§4–§8)	Independent probe	Agreement
Additive side homology	path graph, $H_1 = 0$	$\mathbf{b} = (1, 0, 0)$ for $N = 10, \dots, 10^4$	✓
Connection-rule sensitivity	combinatorial $c_1 = V - E$ varies	$\mathbf{b}_1 \in [0, 795]$ over six rules ($N = 300$)	✓
Pairing identity	relative residual 1.64×10^{-2} at $x = 10^2$	relative residual $\leq 1.64\%$, decaying	✓
Spectral rigidity	self-adjoint spectrum real	$\max \operatorname{Im} \lambda = 0$ (self-adjoint), 1.2692 (non-symmetric perturbation)	✓
Multiplicity control	—	Bonferroni threshold $2.3940 > z_{\max} = 2.3876$	(not used here)

The cross-check matters for two reasons. First, it removes the possibility that a numerical statement is an artefact of one particular implementation: the two codes share no routines, and they agree to the stated tolerances. Second, it exposes the one place where the two analyses deliberately differ in emphasis. The probe also reports a multiple-comparison correction on a three-phase null model, showing that the single "significant" phase ($z = 2.3876$) is *not* significant after Bonferroni correction ($z_{\text{crit}} = 2.3940$ for three tests). That result concerns a different (spectral-statistical) question and is not used in the present structural argument, but it is recorded here for completeness, because it illustrates the same discipline that governs §9.1: a numerical "signal" that does not survive proper control is not a signal.

9.4 Reproducibility

All numerical results in §4–§9 are produced by two single-file scripts with no external data dependencies beyond the standard scientific Python stack (§11.1). The self-tests T1–T11 and the probe's eight checks run in a few seconds; the figure generation runs in under a minute. Every number quoted in §4–§9 appears in a machine-readable results file emitted by the same run that produced the figures, so that the text and the data cannot drift apart silently.

Result 9.3 (reproducibility). Running the five scripts from a clean checkout reproduces: the main self-test at 11/11, the gate logic chain (two premises and six gates; seven links closed, one open), the independent probe self-test at 8/8, the baseline/seam self-test at 6/6, the operator-route self-test at 7/7, the dynamics self-test at 9/9, the figures, and the results file containing the numbers of Results 4.1–4.4, 5.3–5.5, 6.1–6.6, 7.1–7.3 and 8.1–8.2. No step requires network access or a prior run.

9.5 Independent verification of the baseline and the seam

The baseline over-determination of §7.7 and the seam lift of §7.8 are supplied by a third, independent single-file module (`'zeta_balance_probe'`), whose six internal checks are collected in Result 9.4. The module shares no code with the main construction: it has its own sieve, its own Legendre symbol, its own high-precision ξ evaluation (`'mpmath'`, 25 digits), and it computes the Gauss sums from scratch. Because the arithmetic route R2 (Proposition 7.12) uses no analytic input, the module also serves as an arithmetic cross-check on the baseline: the value $1/2$ is returned by a computation that never evaluates ζ , Γ , or a functional equation.

Result 9.4 (baseline/seam self-test). The module `'zeta_balance_probe'` passes 6/6:

Test	Content	Result
T1	geometric axis: $\text{argmin}_b D(b) = 1/2$ (defect 5.56×10^{-18})	pass
T2	Gauss-sum self-normalisation: $\sigma^* = 0.50000000$ (worst deviation 4.6×10^{-14})	pass
T3	exponent involution fixed point = $1/2$ (defect 0)	pass
T4	seam lift: unit modulus at $b = 1/2$ (worst 0.0)	pass
T5	baseline contrast: deviation $\geq 6.2 \times 10^{-2}$ at $b \neq 1/2$	pass
T6	three routes agree to spread = 0.0000	pass

Result 9.5 (execution on the field service). The baseline module is registered as a callable tool of the field service (the cognitive runtime described in §11) and was executed there as well as standalone. The service returned the identical readings (6/6 on the self-test, and the same routes/spread on the full report), so the baseline over-determination and the seam lift are reproduced in an execution environment independent of the one used to prepare the manuscript. This is recorded because it removes the possibility that the numbers of §7.7–§7.8 are an artefact of a single interpreter configuration.

The three modules together — the main construction (11/11), the independent structural probe (8/8), and the baseline/seam module (6/6) — exercise every claim of §4–§8 exactly once, with disjoint prime sieves, disjoint zero tables (where a zero table is used at all), and disjoint determination paths.

9.6 Independent verification of the operator route

The operator route of §6.6 is supplied by a fourth, independent single-file module (``zeta_operator_probe``), whose seven internal checks are collected in Result 9.6. The module shares no code with the other three: it does not use the main construction's sieve, it does not use the baseline module's high-precision ξ , and it obtains its Frobenius eigenvalues from its own brute-force point counts.

Result 9.6 (operator-route self-test). The module ``zeta_operator_probe`` passes 7/7:

Test	Content	Result
T1	function field: $ \alpha = \sqrt{p}$ for 123 curves (worst deviation 0.0)	pass
T2	Hasse bound $ a_q \leq 2\sqrt{q}$ for 123/123 curves	pass
T3	function-field zeros: $\text{Re } s = 1/2$ for 123/123 (worst 1.1×10^{-16})	pass
T4	Riemann zero pipeline anchor: first 30 ordinates on $\text{Re } s = 1/2$	pass
T5	density lock: $ N(T) - RvM \text{ main} = S(T) < 1$ at $T = 50$ ($S = 0.5772$)	pass
T6	Selberg (Weyl, r^2) and ζ density laws are structurally distinct	pass
T7	decision step: model confirmed, G6 localised to archimedean positivity	pass

Result 9.7 (execution on the field service). As with the baseline module (§9.5), the operator-route module is registered as a callable tool of the field service and was executed there. The service returned 7/7 on the self-test and the same function-field readings (123/123 curves on the line, worst deviation 1.1×10^{-16}), and reported the arithmetic row as ``Open = G6 (localised)``. The localisation of §6.6 is therefore reproduced in an execution environment independent of the one used to prepare the manuscript.

The four modules together — the main construction (11/11), the independent structural probe (8/8), the baseline/seam module (6/6), and the operator-route module (7/7) — exercise every claim of §4–§8 and the localisation of §6.6 exactly once, with disjoint prime sieves, disjoint zero tables (where a zero table is used at all), and disjoint determination paths.

9.7 Independent verification of the dynamics route

The dynamics route of §6.7 is supplied by a fifth, independent single-file module (``zeta_dissipation_probe``), whose nine internal checks are collected in Result 9.8. The module shares no code with the other four: it builds the super-kernel and the heat flow from the published series [15, 16], evaluates ξ at its own precision, obtains its zero ordinates from an independent root-finder, and integrates the log-gas flow with its own stepper.

Result 9.8 (dynamics self-test). The module ``zeta_dissipation_probe`` passes 9/9:

Test	Content	Result
T1	super-kernel normalisation $\int_0^\infty \text{Phidu} = 1/8\xi(1/2)$ (ratio 1.0000000000)	pass
T2	$H_0(z) = 1/8\xi(1/2 + iz/2)$ (worst relative error 6.7×10^{-16})	pass
T3	first real zero of H_0 at $2\gamma_1$	pass
T4	real-zero count $N(t)$ monotone non-decreasing; defect 0 at $t = 0$	pass
T5	defect > 0 for $t < 0$ ($N = 1$ at $t = -20$; $N = 3$ at $t = -3$)	pass
T6	unfolded spacings are not equal-spacing ($\text{Var}(s) = 0.1031$)	pass
T7	log-gas gradient flow: H decreases monotonically to the equal-spacing minimum ($-36.770 \rightarrow -38.137$, the exact theoretical value)	pass
T8	$G_6 \rightarrow \Lambda = 0$ with bracket $[0, 0.22]$	pass
T9	reciprocal efficiency $\eta \in (0,1)$ and < 1 for ζ ($\eta = 0.772$; $M_{-1} = 0.873 < M_1 = 0.983$)	pass

Test T9 is the numerical confirmation of the formalisation of §6.8: Proposition 6.34 supplies the inequality $\eta \in (0,1]$, Lemma 6.32 the density floor $E \geq (N-1)\rho^2$, and Lemma 6.33 the arithmetic–harmonic deficit $M_1 - M_{-1} > 0$ off the lattice. The check that the *arithmetic* zeros give $\eta < 1$ (rather than $\eta = 1$) is the same numerical fact, read through the efficiency, that makes $\Lambda \geq 0$ [15].

Result 9.9 (execution on the field service). As with the baseline and operator modules (§9.5, §9.6), the dynamics module is registered as a callable tool of the field service (§11) and was executed there as well as standalone. The service returned 9/9 on the self-test and the identical readings ($\eta = 0.772$, bracket $[0, 0.22]$, decision ``equilibrium_forced_attainment_open``). The localisation of §6.7 is therefore reproduced in an execution environment independent of the one used to prepare the manuscript.

The five modules together — the main construction (11/11), the structural probe (8/8), the baseline/seam module (6/6), the operator-route module (7/7), and the dynamics module (9/9) — exercise every claim of §4–§8 and the two localisations of §6.6–§6.7 exactly once, with disjoint prime sieves, disjoint zero tables (where a zero table is used at all), and disjoint determination paths.

10 Discussion

10.1 The three-level mapping

The constructions of this paper arrange themselves into three levels, and we record the arrangement explicitly because it is the paper's main structural claim: the *complex* level, where the two sides are not equivalent (rank gap); the *pairing* level, where they are identical (explicit formula); and the *spectral* level, where the RH content is carried by a unique, density-locked operator.

Structural object	Level	Role
Chain map φ	Complex	$\mathbb{N}$ additive (rank 1) vs multiplicative (rank ∞): value equality, not isomorphism
Rank gap	Complex	The precise reason homotopy equivalence is impossible
Pairing identity	Functional	The two sides are identical as functionals (explicit formula)
δ_ζ (cumulative)	Functional	Well-defined obstruction, not a divergent series
Self-adjoint H	Spectral	Carries the RH content; unique given its spectrum
Li / Weil positivity	Spectral	The computable form of the remaining existence statement
Least-dissipation flow	Spectral (time)	Forces the existence of a baseline as the flow's equilibrium; that baseline is the one fixed in §7.7; reduces the open item to one scalar Λ
Chern character ch	Cohomological	Common cohomology ring; topological dual of the explicit formula

The spectral level carries a second face besides its existence statement: the *time* face of §6.7, in which the least-dissipation flow necessarily possesses a baseline, and that baseline coincides with the symmetry output of §7.7 — it is *met*, not re-derived, because a flow in the real variable cannot select a vertical line — so that the open item is reduced to a single real constant. The lesson is a level correction. The naive search for a topological equivalence of the two sides fails for a theorem-stated reason (§4.3); the correct statement lives one level down (pairings, §5) and one level up (spectra, §6); and the two bridges (§5, §8) connect these levels. The correction is not stylistic: each level carries a different notion of "same", and the paper's contribution is to say at which level the two sides are same and at which they are not.

10.2 Correspondence with the Hilbert–Pólya programme

Hilbert–Pólya element	This paper's result	Correspondence
Operator with zero spectrum	Uniqueness given the spectrum (Theorem 6.1)	Same object; uniqueness half settled
Self-adjointness	Load-bearing (Proposition 6.3)	Essential, not cosmetic
Spectral density	Locked to Riemann–von Mangoldt (Proposition 6.6)	No free parameter
Positivity	Li and Weil criteria (numerical, §9.2)	Existence has a checkable form

Existence	Open (ES), localised (§6.6) and bracketed to one scalar (§6.7)	The single remaining item, now with a quantitative target
-----------	--	---

The comparison's conclusion is that the Hilbert–Pólya programme [4, 17, 18] and the present structural construction are two views of one object — the former from the operator side, the latter from the two-sides and chain-map side — and that they meet exactly at uniqueness and positivity. From the operator side, one begins with a self-adjoint operator and asks what its spectrum must be; from the structural side, one begins with the two presentations and asks what object carries their mismatch. Both routes arrive at the same three statements: the spectrum determines the operator up to unitary equivalence, the density is locked, and the only open item is existence — which the present route further reduces to the single bracketed constant Λ of the de Bruijn–Newman flow (§6.7).

10.3 Relation to trace formulas

The explicit formula of §5 is an instance of the general paradigm "geometric quantity = main term + structured oscillation" shared by the Selberg [3], Weyl, Gutzwiller and Hadamard trace formulas. In that language, the main term x is the convergence target, the oscillation $-\sum_p x^\rho/\rho$ is the structured tail whose discrete source is the zeros, and the tail never vanishes because there are infinitely many zeros. The Chern-character bridge of §8 is the topological counterpart of this duality: it carries the K-theory class of each side to a common cohomology ring, exactly as the explicit formula carries the function of each side to a common value. The two bridges — analytic and topological — are therefore the same duality seen twice.

This also clarifies the relation of the present work to Connes' noncommutative-geometric treatment [4, 14, 29]. There, the explicit formula becomes a trace formula on the noncommutative space of adèle classes, and RH is reduced to the validity of the trace formula, i.e. to the positivity of the Weil distribution. The present structural route reaches the same positivity from the opposite direction: instead of building the space first and deriving positivity as a trace statement, we begin with the rank gap between the two presentations and show that the mismatch cannot be resolved at the level of complexes, which forces the positivity question to appear at the spectral level. The two routes are complementary, and their agreement on the location of the difficulty (positivity, at the spectral level) is itself a consistency check.

10.4 Limitations

The paper's claims are limited in the following explicit ways.

- The **numerical claims are sample-limited**: 20 zero ordinates; $x \leq 10^5$; $N \leq 300$. Larger samples would lower the truncation floor but cannot remove it, since there are infinitely many zeros, and the conditionally convergent sums are always evaluated with a truncation convention.
- The **analytic statements** (chain-map properties, pairing identity, spectral uniqueness, density lock, functional equation, discrete Gauss–Bonnet) are theorems or propositions and hold without sample limitation. The existence statement ES is open.
- The **Chern character is a homotopy invariant** and does not by itself decide RH; §8.3 states this and it is not softened.
- The **Li partial sums** are computed on a finite range; the omitted tail contributes, so the positivity statement is verified on the range, not proved. The criterion itself is quoted as equivalent to RH from [11, 21].
- **Weil positivity** is quoted as an equivalent criterion [12, 13, 14] rather than independently reproved here; the numerical work is confined to Li's criterion.

- The **dynamical reduction** of §6.7 reduces the open item to the single constant Λ of the de Bruijn–Newman flow, with the unconditional bracket $0 \leq \Lambda \leq 0.22$ taken from [15, 16]. The bracket is a *containment*, not a determination: it does not decide the sign question, and the least-dissipation argument forces the equilibrium rather than the attainment (§6.7.3). The dynamics readings ($\eta = 0.772$, $\text{Var}(s) = 0.1031$) are computed on the first 40 ordinates and inherit the same sample limitation as the rest of the numerical layer.
- The paper **does not prove RH**, does not claim to, and the structural route followed here terminates at ES (§6.5), which §6.6 localises and §6.7 brackets. Any reading according to which the uniqueness statement "proves uniqueness of RH" must be corrected: the uniqueness is of the operator *given its spectrum*, and RH is exactly the statement that a self-adjoint operator with that spectrum exists.

10.5 Falsifiability

The construction is falsifiable at each level, and we record the falsifiers explicitly.

- If a non-trivial zero is found off the critical line, RH fails, ES fails, and all conclusions conditional on RH collapse. The constructions of §4, §5, §7 and §8 would survive, since they do not assume RH; the spectral reading of §6 would fail at the existence gate.
- If λ_n becomes negative for some n , Li's criterion refutes RH, and the positivity form of ES is false.
- If the functional equation fails at higher precision than the tested 10^{-19} relative level, the analytic-continuation correspondence of §7 is refuted. (Given that the equation is a theorem, what such a failure would indicate is an implementation error, which is exactly what the check is designed to detect.)
- If the pairing identity of §5 fails beyond the truncation floor in a way that grows with x , the identification of the two sides as pairings is refuted.
- If a same-spectrum, non-self-adjoint operator were exhibited that *is* unitarily equivalent to a self-adjoint one, Proposition 6.3 would be refuted. (It cannot be, by the invariant $\dim \ker(T - \lambda I)^k$, but the falsifier is stated for completeness.)
- In the dynamical reading, RH is the assertion $\Lambda = 0$. This is falsifiable from *below*: any certified improvement of the Rodgers–Tao bound [15] to some $\Lambda \geq c > 0$ would refute RH. It is settled from *above*: a certified improvement of the Polymath 15 bound [16] to $\Lambda \leq 0$ would prove RH. The present bracket $[0, 0.22]$ decides neither, and the numerical energy reading ($\eta = 0.772 < 1$, the arithmetic zeros sitting strictly above the lattice ground state) is exactly the input that makes the lower end 0 and not below.

10.6 Relation to the probability-cloud reading of the zeros

A companion analysis [30] reads the zeros through a "probability-cloud" framework in which a claim about an infinite domain is replaced by a structured limiting distribution with two layers: an *intrinsic* layer (how far the object is from its target) and a *method* layer (how far any finite method is from resolving it). The present construction is consistent with that reading and sharpens it in three places.

First, the structured tail of the cloud — the oscillating term — is exactly the zero side of our pairing, $-\sum \rho x^{\wedge p/p}$; its discrete source is the zero set, and it never vanishes because there are infinitely many zeros. Second, the "method layer" corresponds to the truncation floor observed in §5.5 and §7.5: the relative residual of 1.64×10^{-2} at $x = 10^2$ is a method artefact that contracts with the number of zeros, not an intrinsic deviation. Third, and most importantly, the present analysis supplies what the cloud picture leaves implicit: *why* the two layers differ. The intrinsic layer is governed by the pairing identity (a theorem), while the method layer is governed by the truncation of a conditionally convergent sum; the two are structurally distinct, and only the first is RH-

relevant. This is precisely the distinction, made in §5.4, between the identity " $\delta_\zeta = 0$ " and the positivity condition that actually carries the RH content.

The comparison therefore runs in both directions: the cloud framework gives an intuitive morphology (intrinsic versus method layer), and the present construction gives the arithmetic reason for that morphology.

10.7 What the construction does not do

For clarity, we state what the paper does **not** accomplish, so that no reader mistakes the reduction for a proof.

- It does not construct the operator H of ES. It proves that if such an operator exists, it is unique and density-locked, and that the question of its existence is equivalent to a computable positivity condition — and, through the dynamical route, to the value of the single constant Λ of the de Bruijn–Newman flow, known to lie in $[0, 0.22]$ with RH exactly $\Lambda = 0$.
- It does not establish RH. It relocates the RH content from a numerical coincidence (two presentations agreeing) to a structural and positive statement (a self-adjoint operator with a specified spectrum), reduces the remaining existence question to one bracketed real number, and shows that the naive relocation (a chain map with vanishing obstruction) is impossible.
- It does not claim that the Chern character, or any homotopy invariant, can decide RH. The bridges are comparisons, not criteria.
- It does not replace the classical programme; it separates its two halves and settles one of them completely.

10.8 Summary of the corrected reading

The following table records, side by side, the naive reading of each structural item and the corrected reading established in this paper.

Item	Naive reading	Corrected reading	Where
Two sides	Homotopy-equivalent complexes	One pairing, written twice	§4.3, §5.3
Obstruction δ_ζ	Pointwise series, diverges on the critical line $\rightarrow$ undefined	Cumulative functional, well-defined for all $x > 1$	§5.1, §5.2
$\delta_\zeta = 0$	Criterion equivalent to RH	Unconditional identity	§5.4
Hilbert–Pólya	"Find the operator" (single problem)	Uniqueness (theorem) + existence (positivity)	§6.2, §6.4
Self-adjointness	A technical hypothesis	Load-bearing	§6.2
Existence	An undefined notion / honesty gap	A computable positivity problem (Li, Weil), localised to archimedean positivity (§6.6) and bracketed to the scalar $\Lambda = 0$ (§6.7)	§6.4, §6.6, §6.7
Continuation	Two sides defined in different regions $\rightarrow$ gap	Unique continuation; critical line is the fixed line	§7
Critical line	A convention, or an input smuggled in through the functional equation	A structural output: unique reflection axis (analytic) and Gauss-sum self-normalisation exponent (arithmetic), which coincide	§7.7

The line as a <i>time</i> limit	Not addressed (the line taken as static geometry)	The flow forces an equilibrium, which coincides with the baseline fixed in §7.7; RH becomes $\Lambda = 0$	§6.7
Discrete vs continuous	Objects of different type ("hard-wired continuity")	One object across a seam: a normalisation, pinned at 1/2	§7.8
Common language	Sought at the level of complexes	Supplied by two bridges (value and class)	§5, §8

10.9 Relation to other equivalent criteria and to the literature

The reduction of §6.5 isolates existence (ES) as the single open item and gives it two computable equivalent forms, Li's criterion [11, 21] and Weil positivity [12, 13, 14]. It is worth locating this within the wider set of equivalent formulations, both to indicate what the structural route shares with them and to indicate what it does not do.

A large number of conditions are known to be equivalent to RH, of which the following are the ones most closely related to the present construction.

- **Explicit-formula criteria.** RH is equivalent to a bound on the error term in the prime number theorem of the form $\psi(x) = x + O(x^{1/2} \log^2 x)$ (the von Mangoldt form of the equivalence). The pairing identity (5.5) is exactly the exact form of this error term, written as $-\sum \rho x^{\rho/p}$; the criterion is what one obtains by bounding that sum. In this language, the present §5 establishes the *identity* underlying the criterion, and §6.4 records the *positivity* criterion that replaces the crude bounding.
- **Positivity criteria.** Li's criterion [11] and Weil's criterion [12, 13] are the two standard positivity conditions; the generalisation due to Bombieri and Lagarias [21] shows that Li's criterion follows from Weil's, and that both apply to arbitrary multisets satisfying the functional equation. The present paper uses Li's criterion as the numerically accessible form and records that the two are equivalent, rather than reproving the equivalence.
- **Spectral criteria.** The Hilbert–Pólya programme [4, 17, 18] and its noncommutative-geometric realisation [4, 14, 29] are spectral in the same sense as §6: they seek an operator whose spectrum is the zero set, or a trace formula whose validity encodes RH. The difference is one of direction. Those programmes build the space (a quantised dynamical system, or the adèle class space) and then derive positivity as a trace statement; the present route begins with the rank gap between the two presentations and shows that the mismatch cannot be resolved at the level of complexes, which forces the positivity question to appear at the spectral level. The two routes meet at the same positivity, and their agreement on the location of the difficulty is a consistency check rather than a competition.
- **Heat-flow / dynamical criteria.** The de Bruijn–Newman flow [19, 20] turns RH into the statement $\Lambda \leq 0$, and with the two unconditional bounds $\Lambda \geq 0$ [15] and $\Lambda \leq 0.22$ [16] into $\Lambda = 0$ with $0 \leq \Lambda \leq 0.22$. This is the criterion §6.7 uses to convert the open item into a bracketed scalar; the least-dissipation identity of [15] supplies the mechanism (relaxation of the zeros to the equal-spacing equilibrium) and the numerical energy reading of Result 6.6. It is the one criterion in this list that carries a *quantity* rather than only a form of the open statement.
- **Statistical criteria.** Montgomery's pair-correlation conjecture [5] and its numerical support [17] describe the *statistics* of the zeros. These are statements about the distribution of the ordinates given that the zeros lie on the critical line; they do not address the structural provenance of the line itself, and they are

therefore orthogonal to the present construction. This is why the multiple-comparison result reported in §9.3 is recorded as "not used here": it belongs to the statistical question.

What the structural route does **not** do, stated plainly: it does not provide an independent proof of any of the classical equivalences (it quotes them); it does not compete numerically with the verification to height 3×10^{12} [6]; and it does not supply the operator of ES. What it does is to *order* these criteria relative to the two-sided construction: it identifies which of them express an identity (§5), which express a theorem given the spectrum (§6.2, §6.3), which express the single open item (§6.4), and which reduce that item to a quantity (§6.7). The value of that ordering is that it removes from the open item everything that can be settled without solving it — uniqueness, density, and the pairing itself — leaving a residual statement of a single, definite type, first localised (§6.6) and then bounded (§6.7).

10.10 On the two bridges as a completeness question

One structural question is left open by the paper and is worth recording, because it is well posed independently of RH.

The two sides are identical at the level of pairings, via the explicit formula (an analytic bridge to a common value, §5), and they admit a second comparison at the level of classes, via the Chern character (a topological bridge to a common ring element, §8). Both bridges realise the same idea — a target that is weaker than the complex level, at which the two sides can be compared without requiring their carriers to be isomorphic. The natural question is whether these two are the only such targets, i.e. whether the class of common targets under which the two sides agree is exhausted by (values, classes) or contains further, structurally distinct members.

This question is well posed because it does not mention RH: it asks about the functorial landscape of §2.9, not about the location of the zeros. A negative answer (there are further bridges) would reveal structure that the present construction has not used; a positive answer (the list is exhausted) would explain why the two-sided tension is resolved at exactly these two levels and no others. Either way the question is checkable in principle, and it is the natural next structural step. We record it here rather than pursuing it, because pursuing it requires machinery — a classification of admissible common targets — that lies outside the scope of the present reductions.

10.11 Ledger of gaps, and their closure

An earlier treatment of this line of argument recorded a list of "gaps" — items it could not close. The purpose of this subsection is to account for every one of them, in one place, with a definite disposition. The guiding principle is strict: an item is marked **closed** only when this paper contains its proof, and the one item that is a genuine open problem is named. Several items that were earlier described as irreducible "limitations" turn out to be *theorems* — statements of non-decidability — and are upgraded accordingly. The upgrade is the content of the following theorem, which closes the entire homotopy-invariant class at once.

Theorem 10.1 (non-decidability of the bridge invariants). Every homotopy invariant of the two-sided complex — the Betti numbers, the Euler characteristic, the Chern character, and every decomposition invariant that is a function of the homotopy type — depends only on the pair of ranks $(1, \pi(N))$. It is therefore the same for ζ and for the Euler product of any other set of "primes" of the same cardinality, and it cannot decide a statement, such as RH, that separates ζ from other Euler products.

Proof. The additive side is a connected graph, so $H_-(C_{\text{add}}) = (\mathbb{Z}, 0, 0, \dots)$ for every additive model — the same multiset whether the model is the path graph or any other connected graph (Result 4.3). The multiplicative side is $\mathbb{Z}^\wedge(\mathbb{P})$ concentrated in degree 0, so $H_-(C_{\text{mul}}) = (\mathbb{Z}^\wedge(\pi(N)), 0, \dots)$, determined by the

cardinality $\pi(N)$ alone. A homotopy invariant is by definition a function of the homotopy type, hence of the homology, hence of the pair $(1, \pi(N))$. This pair is unchanged if the generating set is replaced by any other set of the same size, while the arithmetic object (the Euler product, and with it the zero set) is changed. A quantity constant across objects of the same cardinality cannot distinguish them; RH is a distinguishing statement about one such object; hence no such invariant decides RH. $\square$

Remark 10.2 (what this closes). Theorem 10.1 is not a restatement of "the bridge is only a bridge". It is the proof of it, and it delimits exactly which devices can and cannot carry a decision: the deciding content, if any, must live at the level of *positivity* (§6.4), not at the level of any combinatorial or homotopical invariant of the two sides. This is why the two bridges of §5 and §8 are comparisons and not criteria, and the theorem makes that a matter of proof rather than of caution.

With Theorem 10.1 in hand, the ledger closes as follows.

#	Gap as earlier recorded	Disposition	Where
1	Additive side modelled by a path graph gives $b = (1, 0, 0)$ trivially, carrying no information	Closed — the triviality is the invariance that makes the rank gap robust; the content is the rank gap itself, a theorem	Thm 10.1, §4.5
2	Changing the connection rule makes b_i range over $[0, 795]$; the complex is a construction artefact	Closed — an instance of Theorem 10.1: a homotopy invariant cannot decide RH, so the bridge is non-decisive by proof	Thm 10.1, §8.3
3	The discrete reciprocity symbol ± 1 and the continuous local angle θ_p are objects of different type ("hard-wired continuity")	Closed — the passage is a normalisation, pinned at $1/2$; the seam lift	§7.8, Prop. 7.15
4	The explicit formula is itself a proved bridge, so the "primes $\leftrightarrow$ zeros" seam is not undefined	Closed — the pairing identity is a theorem; the obstruction is a well-defined distribution	§5.3, §5.7
5	"The zeros are trapped to the line" is just a reduction to Hilbert–Pólya, with no operator constructed	Localised twice, and quantified — (i) the operator route is exhibited and its theorem proved in two second structures (function fields, §6.6.1; compact surfaces, §6.6.2), localising G_6 to archimedean positive-definiteness (Weil positivity); (ii) dynamically, least dissipation forces the equal-spacing equilibrium, which <i>is</i> the baseline $1/2$ already fixed in §7.7 — the flow cannot select the line itself, so the value is met, not derived — and the remainder is reduced to the single scalar $\Lambda = 0$, unconditionally bracketed $0 \leq \Lambda \leq 0.22$ (§6.7). It remains	§6.4, §6.6, §6.7

		open in the arithmetic case, but is no longer unlocated, and now has a quantitative two-sided form	
6	The one phase that survived a naive test is no longer significant after multiple-comparison correction	Closed — a definite numerical result (Bonferroni); the item is settled, not deferred	§9.3
7	The chain complex used matches the source manuscript only bespoke-ly	Closed — the chain map is canonical (forced by unique factorisation) and the rank-gap theorem is derived, not posited	§4.1, §4.3
8	The second step of the programme is not well posed: the two sides have no common base space / chain map	Closed — uniqueness of continuation supplies the common base; the chain map exists and is injective	§7.2, §4.2
9	Equality of Betti numbers would not imply RH (homotopy invariants are too coarse)	Closed — Theorem 10.1: no homotopy invariant decides RH; the paper claims no Betti criterion	Thm 10.1
10	The source manuscript is internally inconsistent (an "empirical promise" versus a "mathematical proposition")	Closed — this construction is stated throughout as a mathematical argument and does not inherit the inconsistency	§10.4, §12
11	The de Rham / Dehn side is not calibrated against ζ	Closed — a Dehn-type invariant is homotopy-invariant data of the complex and falls under Theorem 10.1	Thm 10.1
12	Uniqueness (the second half of Hilbert–Pólya)	Closed — the uniqueness half is the spectral theorem; existence is ES	§6.2, §6.6

After the ledger, the accounting is complete: **eleven of the twelve items are closed by a named theorem or proposition of this paper**, and exactly one — item 5, the existence statement ES — remains open. That single item is not left as a bare hole: §6.6 **localises** it to one named ingredient, positive-definiteness at the archimedean place, an object that is automatic in two directly comparable structures and conditional here; and §6.7 **reduces and brackets** it further, to the value of a single real constant, the de Bruijn–Newman parameter Λ , known unconditionally to lie in $[0, 0.22]$ with RH exactly the assertion $\Lambda = 0$. Two points are worth stating plainly, because the earlier treatment understated them.

First, items 1, 2, 9 and 11 were earlier described as *limitations* of the construction. They are in fact **theorems**: a robustness statement (item 1) and three instances of presentation non-decidability (Theorem 10.1). Recording them as limitations, rather than as proofs, understated the result — an invariant that provably cannot decide RH is a stronger statement than an invariant that merely has not decided it. The upgrade is the second of the paper's structural corrections at the level of the bridges.

Second, item 5 is **not** softened but it is **sharpened**, twice. ES is the existence statement of Hilbert–Pólya; it implies RH; and it is not settled here. What §6.6 adds is a localisation, not a resolution: the open item is shown

to be equivalent to the positivity of a single explicit quadratic form at the archimedean place — an object that exists and is positive in the two second structures where the analogue of RH is a theorem. What §6.7 adds is a *quantity*: through the de Bruijn–Newman flow the same item becomes the value of one real constant, $\Lambda = 0$, with an unconditional two-sided bracket $0 \leq \Lambda \leq 0.22$ and a mechanism attached (least dissipation forces the equal-spacing equilibrium — whose location is the baseline already fixed in §7.7, not an output of the flow — and RH is the further statement that the arithmetic datum is already there). The residue is therefore no longer "a quadratic form of unknown sign" but "one number in a known interval". Everything else the route touches is either proved or has been shown to be undecidable by proof — and the distinction between "proved", "proved undecidable", and "open but localised and bracketed" is exactly what a heterogeneous list of complaints lacks.

11 Methods

This section documents the implementation. The **core** of the implementation is not the numerical assembly and not the plotting; it is the **logic chain** — the code that assembles the constructions of §4–§8 into one ordered chain of implications and decides, gate by gate, which links are closed and which is not. Everything else (the arithmetic routines, the spectral routines, the plotting) exists to feed that chain. We therefore present the chain first (§11.2), then the constructions that supply its inputs (§11.3–§11.6), then the data, the tests and the honesty statement. The complete code is reproduced in Appendix A.

11.1 Environment

- Python ≥ 3.11 .
- ``numpy`` (arrays, linear algebra, spectra).
- ``mpmath`` with ``mp.dps = 30`` (high-precision ξ evaluation for the functional-equation check).
- ``matplotlib`` is used **only** for figure rendering, is imported lazily inside a dedicated function, and is not on the path of any determination reported in this paper (§11.7).

The module is single-file with no external data dependencies. The only tabulated external input is the list of the first 20 non-trivial zero ordinates (§11.6). The primes, the von Mangoldt values, the Betti numbers, the spectra and the graphs are all generated at run time. This design ensures that a mistake in the zero table or the prime table would appear as an inconsistency between the two sides of the pairing identity rather than as a silently shared error.

11.2 Core code: the logic chain

The paper's argument is a chain of implications, and the implementation mirrors it: one function per link, each returning a *decision* rather than merely a number, assembled by a driver that returns the gate table and the residual statement. This is the code that corresponds to the reasoning of §4–§6; it is deliberately separated from the numerical routines so that the dependency of each gate on its evidence is explicit.

The chain has **two premises and six gates**. The premises (P1, the structural baseline, and P2, the seam lift) are computed by a separate, independent module and are presented in §11.11; they fix the line on which the gates are stated. The six gates are computed by the driver of §11.2 and are the reduction of Theorem 6.10.

Gate	Claim	Input readouts	Establishes
G1	φ is a homomorphism, injective, not surjective	A (recomputed)	Theorem 4.5
G2	rank 1 vs N $\Rightarrow$ no homology	A ranks	Theorem 4.5

	isomorphism		
G3	prime side = zero side (pairing identity)	B	Theorem 5.5
G4	same-spectrum self-adjoint $\Rightarrow$ unitary equivalence; self-adjointness load-bearing	D	Theorem 6.1 + Proposition 6.3
G5	eigenvalue count = Riemann–von Mangoldt count	D (density)	Proposition 6.6
G6	ES: there exists self-adjoint H with $\text{spec}(H) = \{\gamma_n\}$	open (localised §6.6, bracketed §6.7)	<i>not</i> established — localised and reduced to a bracketed scalar

The first function below states the implication on which the whole programme depends; the second computes the gate table from the readouts of §11.3–§11.5, so that no gate status is hard-coded.

```
def existence_implies_rh() -> Dict:
    """Existence  $\Rightarrow$  RH (pure logic chain, no numerical input)."""
    return {
        'statement': 'If there is a self-adjoint  $H$  with  $\text{spec}(H) = \{\gamma_n\}$ , then RH holds.',
        'proof_chain': [
            '(i)  $H$  self-adjoint  $\Rightarrow \text{spec}(H)$  is real (spectral theorem).',
            '(ii) Hypothesis gives  $\text{spec}(H) = \{\gamma_n\}$ , hence every  $\gamma_n$  is real.',
            '(iii) Each non-trivial zero is  $\rho_n = 1/2 + i\gamma_n$ ;  $\gamma_n$  real  $\Rightarrow \text{Re}(\rho_n) = 1/2$ .',
            ' $\Rightarrow$  All non-trivial zeros lie on the critical line, i.e. RH.',
        ],
        'hypothesis_needed': 'existence (Hilbert-Polya) -- open.',
        'status': 'conditional (existence  $\Rightarrow$  RH); the antecedent is open',
    }

def reduce_structural_program(N: int = 200, K: int = 20) -> Dict:
    """Assemble the A-F readouts into the six-gate decision table.

    Difference from report(): report() only returns readouts. This function
    DECIDES open/closed for each gate from those readouts and returns the
    residual statement.
    """
    A = chain_map_arithmetic(N)          # link 1 / link 2
    B = delta_zeta_cumulative(K=K)       # link 3
    li = li_coefficients(K=K)           # positivity form of link 6
    D = spectral_uniqueness()            # link 4
    W = weyl_density_consistency(K=K)    # link 5

    # G1: the three properties of the chain map, recomputed from raw data
    g1 = bool(A['homomorphism_ok'] and A['injective']
              and (A['surjective_onto_ZP'] is False))

    # G2: rank gap  $\Rightarrow$  no homology isomorphism (derived from A's rank readout)
    g2 = bool(A['rank_multiplicative_side'] > A['rank_additive_side'])

    # G3: pairing identity (both sides agree within the truncation floor)
    worst_rel = max(r['rel_to_x'] for r in B['rows'])
    g3 = bool(worst_rel < 0.05)

    # G4: uniqueness AND load-bearing self-adjointness
    g4 = bool(D['spectra_equal'] and D['matrices_differ']
              and D['jordan_same_spectrum'] and (D['jordan_is_selfadjoint'] is False)
              and (D['unitary_equiv_selfadjoint_vs_jordan'] is False))
```

```

# G5: density lock
g5 = bool(W['consistent'])

# G6: existence -- stays open regardless of the other gates
g6 = False

gates = [
  {'gate': 'G1 chain map', 'claim': 'phi: homomorphism, injective, not surjective',
   'closed': g1, 'source': 'A_chain_map (recomputed)', 'establishes': 'Theorem 4.5'},
  {'gate': 'G2 two sides', 'claim': 'rank 1 vs N => no homology isomorphism; homotopy equivalence impossible',
   'closed': g2, 'source': 'A_chain_map ranks', 'establishes': 'Theorem 4.5'},
  {'gate': 'G3 pairing', 'claim': 'prime side = zero side (worst relative residual < 0.05)',
   'closed': g3, 'source': 'B_delta_zeta_cumulative', 'establishes': 'Theorem 5.5'},
  {'gate': 'G4 uniqueness', 'claim': 'same-spectrum self-adjoint => unitary equivalence; self-adjointness load-bearing',
   'closed': g4, 'source': 'D_spectral_uniqueness', 'establishes': 'Theorem 6.1 + Prop 6.3'},
  {'gate': 'G5 density', 'claim': 'eigenvalue count = Riemann-von Mangoldt count',
   'closed': g5, 'source': 'D_weyl_density', 'establishes': 'Proposition 6.6'},
  {'gate': 'G6 existence', 'claim': 'ES: exists self-adjoint H with spec(H) = {gamma_n}',
   'closed': g6, 'source': 'open (equivalent: Li >= 0, Weil >= 0)',
   'establishes': 'not established -- the single remaining item'},
]
n_closed = sum(1 for g in gates if g['closed'])
return {
  'gates': gates,
  'n_closed': n_closed,
  'n_open': len(gates) - n_closed,
  'residual_statement': 'exists self-adjoint H with spec(H) = {gamma_n}    (ES = Hilbert-Polya)',
  'positivity_check': {
    'li_all_positive_n_le_20': li['all_positive'],
    'lambda_1_partial': li['lambda_1_partial'],
    'lambda_1_exact': li['lambda_1_exact_analytic'],
  },
  'implies': existence_implies_rh(),
  'verdict': ('Five gates close as theorems/propositions; the chain terminates at G6. '
             'Hence the structural programme reduces to ES, whose equivalent positivity '
             'form (Li) is verified positive on the tested range.'),
}

```

Creation logic. Three decisions determine whether this driver is a genuine logic chain or a set of hard-coded labels.

First, **each gate is decided from its readout, not asserted**. `g1` is a conjunction of the three properties returned by `chain\_map\_arithmetic`, so that removing any one of them (for instance, if `injective` were somehow `False`) would flip the gate. `g2` is derived from the *rank comparison* rather than from the name of the theorem, so that a change in the model that altered the ranks would change the gate. `g3` is decided by a numeric tolerance on the residual. `g4` requires the conjunction of the four spectral facts, so that the load-bearing property cannot report closure unless the Jordan block is simultaneously same-spectrum and non-self-adjoint. `g5` is decided by the density consistency flag.

Second, **G6 is the only gate that is not computed from readouts**, and it is set to `False` unconditionally. This is the honest encoding of the paper's logic: existence is not decided here, and no readout available to this module could decide it. The gate is nevertheless *carried* in the table — with its equivalent positivity form attached, `positivity\_check` — so that the chain reports the residual as an item with computable content rather than as an absence.

Third, **the implication is separated from the decision**. `existence\_implies\_rh()` contains no numbers: it is the logical content of §2.4 and §6.4, namely that a self-adjoint operator with the zero spectrum forces real γ_n and hence RH. Keeping the implication and the gate table separate means that the conditional nature of the result is

visible in the code: the programme returns `n\_closed = 5`, `n\_open = 1`, and the implication `ES  $\Rightarrow$  RH`, and nothing in the code asserts RH.

Running the driver on the readouts gives, deterministically:

G1 chain map	closed=True	phi: homomorphism, injective, not surjective
G2 two sides	closed=True	rank 1 vs N => no homology isomorphism; homotopy equivalence impossible
G3 pairing	closed=True	prime side = zero side (worst relative residual < 0.05)
G4 uniqueness	closed=True	same-spectrum self-adjoint => unitary equivalence; self-adjointness load-bearing
G5 density	closed=True	eigenvalue count = Riemann-von Mangoldt count
G6 existence	closed=False	ES: exists self-adjoint H with spec(H) = {gamma_n}
closed=5 open=1		
residual: exists self-adjoint H with spec(H) = {gamma_n}		(ES = Hilbert-Polya)

This is Theorem 6.10 stated as running code: the structural programme reduces to ES, and to nothing else. The *further* reduction of G6 — from "an operator exists" to "the scalar Λ vanishes", with its bracket — is a separate decision step, computed by the dynamics module of §11.13 and localising (not closing) the gate in the sense of §6.6–§6.7.

11.3 Core code: the chain map (links G1, G2)

The chain map of §4 is built from the factorisation of each integer into prime powers. The factorisation routine is trial division with the standard 2-then-odd optimisation; it returns a dictionary of exponents, which the chain-map routine converts into a coordinate vector in the basis of primes up to N. The three properties of §4.2 are checked directly, and the two ranks are read off.

```
def factor_exponents(n: int) -> Dict[int, int]:
    """Return {p: v_p(n)} for n = product p^{v_p(n)} (trial division)."""
    out: Dict[int, int] = {}
    m, d = n, 2
    while d * d <= m:
        while m % d == 0:
            out[d] = out.get(d, 0) + 1
            m //= d
        d += 1 if d == 2 else 2          # 2, 3, 5, 7, ...
    if m > 1:
        out[m] = out.get(m, 0) + 1
    return out

def chain_map_arithmetic(N: int = 200) -> Dict:
    """Chain map phi given by the fundamental theorem of arithmetic.

    Carrier: the free abelian group Z^{(P)} on the primes <= N.
    Each integer n <= N is sent to the vector of its prime exponents.
    Homomorphism / injective / non-surjective are each recomputed from data.
    """
    ps = [p for p in zg.sieve(N)]
    idx = {p: i for i, p in enumerate(ps)}

    def phi(n: int) -> np.ndarray:
        v = np.zeros(len(ps), dtype=np.int64)
        for p, a in factor_exponents(n).items():
            if p in idx:
                v[idx[p]] = a
        return v

    # -- G1a: homomorphism phi(ab) = phi(a) + phi(b), on sampled pairs ----
    hom_ok, hom_rows = True, []
    for (a, b) in [(6, 35), (12, 91), (17, 19), (60, 77), (128, 125)]:
```

```

    if a * b <= N:
        ok = bool(np.array_equal(phi(a * b), phi(a) + phi(b)))
        hom_ok = hom_ok and ok
        hom_rows.append({'a': a, 'b': b, 'ab': a * b, 'hom_ok': ok})

# -- G1b: injectivity, no collision among n = 1..N -----
seen, injective, first_collision = {}, True, None
for n in range(1, N + 1):
    key = tuple(int(x) for x in phi(n))
    if key in seen:
        injective, first_collision = False, (seen[key], n)
        break
    seen[key] = n

# -- G1c: non-surjectivity, the element -[2] has no preimage -----
target = np.zeros(len(ps), dtype=np.int64)
if len(ps) > 0:
    target[0] = -1 # -[2]
has_preimage = any(np.array_equal(phi(n), target)
                    for n in range(1, N + 1)) if len(ps) > 0 else False

# -- G2: the two ranks. (N,+) is generated by {1}; (N,.) by the primes -
return {
    'N': N,
    'n_primes': len(ps),
    'homomorphism_ok': bool(hom_ok),
    'hom_rows': hom_rows,
    'injective': bool(injective),
    'first_collision': first_collision,
    'surjective_onto_ZP': bool(has_preimage),
    'rank_additive_side': 1,
    'rank_multiplicative_side': len(ps),
    'sample_vectors': {str(n): phi(n).tolist() for n in (1, 2, 6, 12, 36, 210)},
}

```

Creation logic. Three decisions are worth recording. First, `phi` allocates a fresh zero vector per call rather than mutating a shared buffer; this makes the injectivity scan — which stores N vectors — correct without aliasing bugs, at memory cost proportional to N , negligible at $N = 200$ and the first thing to change for large N . Second, injectivity is decided by a dictionary keyed on the exponent tuple, so that a collision would be reported together with the colliding pair, not merely as a boolean. Third, the ranks are returned as first-class fields `rank_additive_side` and `rank_multiplicative_side`, because gate G2 is decided by comparing them; had they been inlined as literals in the driver, the gate would have been a label rather than a decision.

11.4 Core code: the cumulative obstruction (link G3)

The cumulative obstruction is computed on two independent sides and compared at four values of x .

```

def delta_zeta_cumulative(xs=(100.0, 1000.0, 10000.0, 100000.0), K: int = 20) -> Dict:
    """Cumulative obstruction and two-sided agreement.

    delta_zeta(x) := psi(x) - x + log 2pi + (1/2) log(1 - x^{-2})
    By the explicit formula this equals the zero-side expression
    delta_zeta(x) = -sum_rho x^rho / rho.
    Both sides are computed here independently.
    """
    rows = []
    for x in xs:
        L = zg.von_mangoldt_table(int(x)) # primes -> Lambda table
        psi = zg.psi_exact(x, L) # exact prime side
        prime_side = psi - x + math.log(2 * math.pi) + 0.5 * math.log(1.0 - x ** -2.0)
        zero_side = 0.0
        for g in zg._ZETA_ZEROS[:K]: # first K zero ordinates

```

```

        for rho in (complex(0.5, g), complex(0.5, -g)): # conjugate pair
            zero_side += -(x ** rho) / rho.real
        rows.append({'x': x, 'K': K,
                    'delta_prime_side': round(prime_side, 6),
                    'delta_zero_side': round(zero_side, 6),
                    'abs_residual': round(abs(prime_side - zero_side), 6),
                    'rel_to_x': round(abs(prime_side - zero_side) / x, 8)})
    return {'rows': rows,
            'definition': 'delta_zeta(x) = psi(x) - x + log2pi + (1/2)log(1-x^-2) = -sum_rho x^rho/rho'}

```

Creation logic. The essential choice is that the residual is normalised by x (`rel\_to\_x`), not reported as an absolute number: the obstruction's scale grows like $\sqrt{x}$ against a main term of size x , so an absolute residual can rise while the relative approximation improves (Result 5.3). Gate G3 is decided on `rel\_to\_x` for exactly this reason. The conjugate pair $(\rho, \bar{\rho})$ is summed explicitly so that the imaginary parts cancel by construction rather than by numerical accident, which keeps `zero\_side` real to machine precision.

11.5 Core code: positivity, uniqueness, density (links G4, G5, and the positivity form of G6)

Li coefficients — the computable equivalent form of the open gate's positivity condition:

```

def li_coefficients(ns=tuple(range(1, 21)), K: int = 20) -> Dict:
    """Li coefficients lambda_n = sum_rho [1 - (1 - 1/rho)^n]: partial sums and positivity.

    Equivalence theorem (Li 1997): RH <=> lambda_n >= 0 for all n >= 1.
    Analytic reference: lambda_1 = 1 + gamma/2 - (1/2) log(4 pi) ~ 0.0230957...
    """
    gs = list(zg._ZETA_ZEROS[:K])
    rows = []
    for n in ns:
        s = 0.0
        for g in gs:
            for rho in (complex(0.5, g), complex(0.5, -g)):
                s += (1.0 - (1.0 - 1.0 / rho) ** n).real
            rows.append({'n': n, 'lambda_partial': round(s, 6)})
    vals = [r['lambda_partial'] for r in rows]
    lam1_exact = 1.0 + 0.5772156649015329 / 2.0 - 0.5 * math.log(4.0 * math.pi)
    return {'K': K, 'rows': rows, 'all_positive': bool(all(v > 0 for v in vals)),
            'lambda_1_partial': rows[0]['lambda_partial'],
            'lambda_1_exact_analytic': round(lam1_exact, 6)}

```

Creation logic. The reference value λ_1 is computed from the closed form (2.8) with the Euler–Mascheroni constant written to 16 digits, so that the comparison of Result 6.4 is a genuine test: the partial sum must fall *below* the analytic value by the omitted tail, and the routine returns both numbers rather than only the sign, so that a modelling error (for example, a wrong conjugate-pair convention) would show up as a discrepancy in the wrong direction. This function supplies the `positivity\_check` field attached to gate G6 in §11.2.

Spectral uniqueness and the load-bearing counterexample (link G4):

```

def spectral_uniqueness(seed: int = 2026, n: int = 6) -> Dict:
    """Same-spectrum self-adjoint operators are unitarily equivalent;
    self-adjointness is load-bearing (Jordan counterexample)."""
    rng = np.random.default_rng(seed)
    A = rng.normal(size=(n, n)); A = (A + A.T) / 2.0 # self-adjoint
    lam = np.sort(np.linalg.eigvalsh(A))
    M = rng.normal(size=(n, n)); Q, _ = np.linalg.qr(M) # random orthogonal
    B = Q.T @ A @ Q
    spec_equal = bool(np.allclose(lam, np.sort(np.linalg.eigvalsh(B)), atol=1e-9))
    matrices_differ = bool(not np.allclose(A, B, atol=1e-9))

```

```

J = np.array([[0.0, 1.0], [0.0, 0.0]]) # Jordan block
jordan_same_spec = bool(np.allclose(np.sort(np.linalg.eigvals(J).real), [0.0, 0.0]))
jordan_selfadj = bool(np.allclose(J, J.T)) # False

g = np.array(zg._ZETA_ZEROS, dtype=float) # zero ordinates
Hz = np.diag(g)
Mz = rng.normal(size=(len(g), len(g))); Qz, _ = np.linalg.qr(Mz)
Hz2 = Qz.T @ Hz @ Qz
zeta_same_spec = bool(np.allclose(np.sort(np.linalg.eigvalsh(Hz2)), np.sort(g), atol=1e-9))
zeta_differ = bool(not np.allclose(Hz, Hz2, atol=1e-9))
return {'n': n, 'spectra_equal': spec_equal, 'matrices_differ': matrices_differ,
        'jordan_same_spectrum': jordan_same_spec, 'jordan_is_selfadjoint': jordan_selfadj,
        'unitary_equiv_selfadjoint_pair': spec_equal,
        'unitary_equiv_selfadjoint_vs_jordan': bool(jordan_selfadj and jordan_same_spec),
        'zeta_spectrum_selfadjoint': bool(np.allclose(Hz2, Hz2.T, atol=1e-9)),
        'zeta_same_spec': zeta_same_spec,
        'zeta_same_spec_but_different_matrix': zeta_differ}

```

Creation logic. The Jordan test is written so that it can only return 'False' for unitary equivalence: the routine computes whether J is self-adjoint (it is not) and *and*-combines that with same-spectrum in 'unitary_equiv_selfadjoint_vs_jordan', so that the load-bearing property is asserted only when both the spectral coincidence and the failure of self-adjointness are present. A naive implementation testing only "spectra agree" would wrongly report equivalence; the 'and' is what encodes Proposition 6.3. Gate G4 is the conjunction of this flag with 'spectra_equal' and 'matrices_differ', so that neither half can be satisfied alone.

Density lock (link G5):

```

def weyl_density_consistency(T: float = 77.1448400689, K: int = 20) -> Dict:
    """Density constraint: the eigenvalue count must equal the
    Riemann-von Mangoldt zero count."""
    main = (T / (2 * math.pi)) * math.log(T / (2 * math.pi * math.e)) + 7.0 / 8.0
    S = K - main
    return {'T': T, 'K': K, 'N_main_rvm': round(main, 6), 'actual_count': K,
            'S_T_residual': round(S, 6),
            'consistent': bool(abs(S) <= 0.35 * math.log(T) + 1.0)}

```

Creation logic. The tolerance $0.35 \log(T) + 1.0$ encodes the classical bound $S(T) = O(\log T)$ with a deliberately loose constant, because the point of the gate is to certify that the residual is of the right order*, not to establish a best constant. A sharp constant here would turn a structural consistency check into a numerical optimisation with no structural content, and would make gate G5 sensitive to a quantity that carries no structural information.

11.6 Data used

The only tabulated external input is the list of the first 20 non-trivial zero ordinates, entered as:

14.134725, 21.022040, 25.010858, 30.424876, 32.935062, 37.586178, 40.918719, 43.327073, 48.005151,
49.773832,
52.970321, 56.446248, 59.347044, 60.831779, 65.112544, 67.079811, 69.546402, 72.067158, 75.704691,
77.144840.

Everything else — the primes, the von Mangoldt values, the Betti numbers, the spectra, the graphs — is generated at run time. The computed values of $\psi(x)$ and of the zero sum are compared against each other rather than against an external table, so that no external numeric data can contaminate both sides of an identity simultaneously.

11.7 What is **not** on the logic path

Two kinds of code exist in the module but are deliberately excluded from the chain.

- **Graph and Betti routines** (``_prime_graph``, ``chern_character_k0``) supply the topological bridge of §8. They are not gates of the chain, because §8.3 establishes that the Chern character is a homotopy invariant and does not decide RH. Reporting them as gates would over-state the bridge; they are therefore reported as a *bridge*, with an explicit non-decision note.
- **Plotting** (``_setup_mpl``, ``make_figures``) renders the figures of the text. It is imported lazily, consumes the same in-memory readouts as the report, and feeds no decision. It is listed in Appendix B as a reproduction step only.

Separating these from the chain is not bookkeeping: it is the difference between the code that *decides* the paper's claims and the code that merely *displays* them.

11.8 Self-tests and reproduction

The module exposes five command-line modes:

```
python zeta_weil_construction.py selftest # runs T1..T11; exits 0 iff 11/11
python zeta_weil_construction.py report   # results JSON, including the gate table
python zeta_weil_construction.py reduce   # prints the gate decision table
python zeta_weil_construction.py figures  # writes the figures (not a decision step)
python zeta_weil_construction.py all      # selftest + report + figures
```

The self-test harness collects each check as a record containing its name, its boolean outcome and a detail string; the summary reports the pass count out of the total, and the process exit code reflects it, so the self-test is usable in continuous integration without parsing output. The eleventh check is the logic chain itself.

```
def selftest() -> Dict:
    checks: List[Dict] = []

    def ck(name: str, ok: bool, detail: str = '') -> None:
        checks.append({'name': name, 'ok': bool(ok), 'detail': detail})

    cm = chain_map_arithmetic(200)
    ck('T1 chain map is a homomorphism', cm['homomorphism_ok'], 'phi(ab)=phi(a)+phi(b)')
    ck('T2 chain map is injective', cm['injective'], 'unique factorisation')
    ck('T3 chain map is not surjective', (cm['surjective_onto_ZP'] is False)
        and cm['rank_multiplicative_side'] > cm['rank_additive_side'],
        'rank %d -> %d' % (cm['rank_additive_side'], cm['rank_multiplicative_side']))

    dc = delta_zeta_cumulative((100.0, 1000.0, 10000.0), K=20)
    worst = max(r['rel_to_x'] for r in dc['rows'])
    ck('T4 cumulative obstruction agrees on both sides', worst < 0.05, 'worst rel=%.6f' % worst)

    li = li_coefficients(tuple(range(1, 16)), K=20)
    ck('T5 Li partial sums all positive', li['all_positive'] is True,
        'lambda_1=%.6f' % li['lambda_1_partial'])

    su = spectral_uniqueness()
    ck('T6 uniqueness: same-spectrum self-adjoint', su['spectra_equal'] and su['matrices_differ'])
    ck('T7 self-adjointness is load-bearing (Jordan)',
        (su['jordan_same_spectrum'] is True) and (su['jordan_is_selfadjoint'] is False)
        and (su['unitary_equiv_selfadjoint_vs_jordan'] is False))

    wd = weyl_density_consistency()
    ck('T8 density consistency (Weyl count)', wd['consistent'], 'S=%.4f' % wd['S_T_residual'])
```

```

ch = chern_character_k0(300, ('near10', 'cong12', 'product'))
ck('T9 Chern character: discrete Gauss-Bonnet holds',
    all(r['gb_ok'] for r in ch['rows']) and all(ch['ring_homomorphism'].values()))

fe = functional_equation_check()
ck('T10 analytic continuation: xi(s)=xi(1-s) at 30 digits',
    bool(fe.get('available')) and bool(fe.get('equation_holds')),
    'worst rel=%2e' % fe.get('worst_rel_diff', float('nan')))

lc = reduce_structural_program(N=120, K=20)
ck('T11 logic chain: five gates close, exactly ES stays open',
    lc['n_closed'] == 5 and lc['n_open'] == 1
    and lc['gates'][-1]['gate'] == 'G6 existence'
    and lc['gates'][-1]['closed'] is False,
    'closed=%d open=%d' % (lc['n_closed'], lc['n_open']))

npass = sum(1 for c in checks if c['ok'])
return {'pass': npass, 'total': len(checks), 'ok': npass == len(checks), 'checks': checks}

```

The suite passes 11/11.

11.9 Verification protocol and independence

The module above is the primary implementation. A second module was written independently to re-derive the same facts from a different angle (§9.3), with its own prime sieve, its own von Mangoldt table, its own zero list and its own graph routines. The two modules share no code. Agreement between them, to the tolerances reported in §9.3, is the basis for the claim that the numerical statements of §4–§9 are not artefacts of a particular implementation. Where the two modules deliberately differ — the multiple-comparison correction on the three-phase null model — the difference is reported rather than suppressed (§9.3).

In total **five single-file modules** supply the paper's claims and share no code: the main construction of §11.2–§11.8 (self-test 11/11), the independent structural probe (§9.3), the baseline-and-seam module of §11.11 (premises P1–P2, 6/6), the operator-route module of §11.12 (localisation of G6, 7/7), and the dynamics module of §11.13 (reduction of G6 to the bracketed scalar, 9/9). Each has its own sieve, its own zero table (where one is used) and its own determination path; the localisations of §6.6 and §6.7 are therefore reproduced across five disjoint implementations.

11.10 Honesty statement

Evidence in this paper is three-layered, and the layer of each claim is stated where the claim is made.

(i) **Analytic / exact** (sample-independent). The chain-map properties (Propositions 4.1–4.3, Theorem 4.5); the pairing identity (Theorem 5.5); the well-posedness of the cumulative obstruction (Lemma 3.5) and the divergence of the pointwise series (Propositions 5.1–5.2); the corrected reading of $\delta_\zeta = 0$ (Proposition 5.8); the spectral uniqueness theorem (Theorem 6.1) and the load-bearing counterexample (Proposition 6.3); the density lock (Proposition 6.6); the reduction theorem (Theorem 6.10); the functional equation and the uniqueness of continuation (Theorem 7.1, Corollary 7.2, Proposition 7.5); the discrete Gauss–Bonnet identity (Proposition 8.2) and the ring-homomorphism property (Proposition 8.3). For the dynamics of §6.7 the exact content is *quoted*: the de Bruijn–Newman existence of Λ and the equivalence $\text{RH} \Leftrightarrow \Lambda \leq 0$ (Theorems 6.22–6.23 [19, 20]) and the least-dissipation identity $\partial_t H = -4E$ with E monotone non-increasing (Theorem 6.24 [15]).

(ii) **Numerical** (sample-limited): the cumulative-obstruction residuals ($x \leq 10^5$, 20 zeros); the Li partial sums ($n \leq 20$); the Weyl-count residual; the Chern-character values ($N = 300$); the functional-equation residuals (30

digits); the dynamics readings — heat-flow zero count and its monotonicity, reciprocal efficiency η , and the log-gas relaxation to the equal-spacing minimum (40 ordinates; §6.7, Result 6.6).

(iii) **Open:** existence (ES = Hilbert–Pólya), with the equivalent positivity forms (Li, Weil) and the equivalent single-scalar form $\Lambda = 0$ with the bracket $0 \leq \Lambda \leq 0.22$ (§6.7). This is gate G6 of the logic chain, and it is the only gate the chain leaves open; §6.6 localises it and §6.7 brackets it, but neither closes it.

Statement. This paper does not prove RH. It proves that, given the spectrum, the operator is unique up to unitary equivalence and has locked density; that the two sides are identical as pairings; that the canonical chain map exists but is not surjective, and that this forbids the naive topological reading; and that the only unresolved item is existence, whose two equivalent positivity forms are numerically positive on the tested range, and whose dynamical form is the value of a single real constant inside a two-sided bracket ($0 \leq \Lambda \leq 0.22$, §6.7). Every item that an earlier treatment recorded as an "honesty gap" is here either closed as a theorem or restated as a sharply defined open problem with a computable equivalent form — and the logic chain of §11.2, together with the two localisations of §11.12–§11.13, is the code that makes that separation explicit.

11.11 Core code: the baseline and the seam (premises P1, P2)

The two premises of §6.5 are computed by a third, independently written module, `'zeta_balance_probe'`, whose role is to fix the line on which the gates are stated. As with the driver of §11.2, the core of this module is **logic-chain code**: each route returns a *decision*, and a driver decides whether each premise closes. The code that *decides* is again separated from the code that merely *displays* — the two panels of Figs. 7–8 are drawn by a plotting routine that is not on the decision path.

The premise driver is the following. Note that P1 is decided as a **conjunction** of two routes with disjoint inputs, which is what makes it a genuine decision rather than a single computation.

```
def baseline_premises(N: int = 600) -> Dict:
    """Premises P1 (baseline) and P2 (seam lift), each decided from its readouts.

    P1 is decided as the CONJUNCTION of two independent routes:
        R1 analytic -- the reflection axis of xi is the minimiser of D(b);
        R2 arithmetic -- the Gauss-sum self-normalisation exponent is 1/2.
    A failure of either route would open P1.
    P2 is decided from the seam-lift residual: unit modulus at b = 1/2 AND its
    failure at every other tested exponent (a one-sided check would accept any b).
    """
    r1 = xi_axis_scan()                # analytic route (mpmath, high precision)
    r2 = gauss_sum_exponent(N)         # arithmetic route (Legendre symbol only)
    sl = seam_lift(N)                  # seam lift

    p1_axis = bool(abs(r1['argmin_b'] - 0.5) <= 0.01 and r1['argmin_defect'] < 1e-12)
    p1_gauss = bool(abs(r2['slope_sigma_star'] - 0.5) < 1e-6)
    p1 = bool(p1_axis and p1_gauss)    # conjunction => a genuine gate
    p2 = bool(sl['worst_dev_from_unity_at_half'] < 1e-9
              and min(a['worst_dev'] for a in sl['baseline_contrast']
                      if a['b'] != 0.5) > 1e-3)
    return {
        'P1_baseline': {'closed': p1, 'r1_axis': r1['argmin_b'],
                        'r2_sigma_star': r2['slope_sigma_star'],
                        'establishes': 'Corollary 7.13'},
        'P2_seam': {'closed': p2,
                    'worst_dev_at_half': sl['worst_dev_from_unity_at_half'],
                    'establishes': 'Proposition 7.15'},
        'n_closed': int(p1) + int(p2),
        'n_open': 2 - (int(p1) + int(p2)),
    }
```

Creation logic. Three decisions make this a logic chain rather than a set of labels. First, **P1 is a conjunction of two routes with disjoint inputs**: R1 evaluates ξ at high precision with ``mpmath``, whereas R2 uses only the Legendre symbol and an exact finite exponential sum and never touches ζ , Γ or a functional equation. The premise therefore certifies *over-determination*: two unrelated structures returning the same constant. Had they disagreed, ``p1_axis` and ``p1_gauss`` would open P1. Second, **P2 is decided together with its contrast** — the same routine that confirms unit modulus at $b = 1/2$ confirms its failure at $b \in \{0.30, 0.40, 0.49, 0.51, 0.60, 0.70\}$; a one-sided check would accept any value of b . Third, **the two routes are returned as separate fields** (``r1_axis``, ``r2_sigma_star``) rather than collapsed into one boolean, so that the agreement is visible as two numbers rather than as a verdict.

The route that carries the *arithmetic* content — and which, because it uses no analytic input, is the strongest single witness for the baseline — is the Gauss-sum route:

```
def gauss_sum_exponent(N: int = 1000, lo: int = 5) -> Dict:
    """R2: the self-normalisation exponent of the quadratic reciprocity character.

    |G(p)| = sqrt(p) => the slope of log|G(p)| against log p is exactly 1/2.
    Pure arithmetic: no zeta, no gamma, no functional equation.
    """
    ps = [p for p in sieve(N) if p >= lo]
    pts = []
    for p in ps:
        re = im = 0.0
        for a in range(1, p):
            c = legendre(a, p)          # (a|p) in {-1, 0, 1}
            if c:
                th = 2.0 * math.pi * a / p
                re += c * math.cos(th); im += c * math.sin(th)
        pts.append((math.log(p), math.log(math.hypot(re, im))))
    n = len(pts)
    sx = sum(x for x, _ in pts); sy = sum(y for _, y in pts)
    sxx = sum(x * x for x, _ in pts); sxy = sum(x * y for x, y in pts)
    slope = (n * sxy - sx * sy) / (n * sxx - sx * sx)
    return {'slope_sigma_star': round(slope, 9), 'n_primes': n,
            'verdict': 'log|G(p)| = sigma*log p with sigma* -> 1/2'}
```

The seam lift decides P2 by the *same* normalisation exponent seen from the other side, and returns its contrast so that the pinning is a decision:

```
def seam_lift(N: int = 1000, lo: int = 5) -> Dict:
    """P2: |G(p)| * p^{-1/2} = 1 for every odd prime; fails at b != 1/2."""
    ps = [p for p in sieve(N) if p >= lo]
    mags = []
    for p in ps:
        re = im = 0.0
        for a in range(1, p):
            c = legendre(a, p)
            if c:
                th = 2.0 * math.pi * a / p
                re += c * math.cos(th); im += c * math.sin(th)
        mags.append((p, math.hypot(re, im)))
    worst = max(abs(m * p ** -0.5 - 1.0) for p, m in mags)          # -> 0
    contrast = [{ 'b': b,
                  'worst_dev': max(abs(m * (p ** -b) - 1.0) for p, m in mags)}
                for b in (0.30, 0.40, 0.49, 0.50, 0.51, 0.60, 0.70)]
    return {'worst_dev_from_unity_at_half': worst, 'baseline_contrast': contrast,
            'verdict': 'unit modulus at b = 1/2 and only there'}
```

Running the premise driver gives, deterministically:

```
P1 baseline    closed=True    R1 argmin_b=0.500    R2 sigma*=0.500000000
```

```
P2 seam lift closed=True worst_dev_at_half=0.0e+00
closed=2 open=0
```

So the two premises close, and the full chain of §6.5 is: two premises closed (§11.11), five gates closed and one gate open (§11.2). The chain as a whole therefore has **seven closed links and one open link**, and the open link is ES. This is the code form of the statement that $1/2$ is a structural output: no branch of the code can return $1/2$ unless the two routes agree, and neither route contains the value $1/2$ as an input.

11.12 Core code: the operator route (the decision step for G6)

Section 6.6 exhibits the operator route in two second structures and localises G6. As in §11.2 and §11.11, the code below is the **decision** step, not a display step: it decides whether the model route is confirmed and reports the status of G6. It is the `'operator_route_check'` entry of the independently written module `'zeta_operator_probe'`.

```
def operator_route_check(T: float = 50.0) -> Dict:
    """Decide whether the operator route is confirmed in the second structures,
    and report the status of G6. The arithmetic row is never returned as closed."""
    ff = function_field_zeros()          # 123 elliptic curves over F_p: eigenvalues alpha
    ri = density_lock(T)                 # arithmetic side: zero positions + Riemann-von Mangoldt

    # ---- decisions (not display) -----
    model_on_line = bool(ff['n_curves'] > 0
                        and ff['n_on_line'] == ff['n_curves']          # every zero on Re s = 1/2
                        and ff['worst_dev_re_from_half'] < 1e-9)
    model_unitary = bool(ff['worst_dev_abs_alpha'] < 1e-9)             # |alpha| = sqrt(q)
    density_ok = bool(ri.get('error') is None
                    and abs(ri['S_T']) < 1.0                          # density lock (link G5)
                    and ri['worst_dev_re_from_half'] < 1e-12)
    decision = 'model_ok' if (model_on_line and model_unitary and density_ok) else 'model_fail'
    # -----
    return {
        'decision': decision,
        'models': {
            'function_field': {'operator': 'Frobenius F on H^1, U = F/sqrt(q)',
                              'selfadjoint_source': 'positive-definite intersection form',
                              'status': 'Theorem (Weil 1948)'},
            'selberg_surface': {'operator': 'Laplace-Beltrami operator Delta_X',
                                'selfadjoint_source': 'positive-definite metric',
                                'status': 'Theorem (Selberg 1956)'},
            'arithmetic': {'operator': 'Hilbert-Polya operator (unnamed)',
                           'selfadjoint_source': 'Weil positivity W(f * ftilde) >= 0',
                           'status': 'Open = G6 (localised, not proved here)'},
        },
        'g6_localisation': 'G6 == existence of a positive-definite pairing at the archimedean '
                           'place (Weil positivity).',
    }
```

The Frobenius eigenvalues are obtained from point counts, not from any table, so the model decision is self-contained:

```
def frobenius_eigs(a, b, p):
    """alpha = (a_p + sqrt(a_p^2 - 4p)) / 2 with a_p = p + 1 - #E(F_p)."""
    N = count_points_elliptic(a, b, p)          # brute force, no table
    ap = p + 1 - N
    abs_alpha = (ap*ap + (4*p - ap*ap))**0.5 / 2.0    # = sqrt(p) exactly, by Hasse
    return {'N_Fp': N, 'a_p': ap, 'abs_alpha': abs_alpha,
            'hasse_ok': abs(ap) <= 2*p**0.5 + 1e-12,
            'dev_abs_alpha': abs(abs_alpha - p**0.5)}
```

and the zero position of each curve is the normalised logarithm of that eigenvalue:

```
s = cmath.log(alpha) / math.log(p)          # the zero position s = log alpha / log p
rows.append({'p': p, 'curve': [a, b], 'abs_alpha': e['abs_alpha'],
            'zero_re': s.real, 'zero_im': s.imag,
            'dev_re_from_half': abs(s.real - 0.5)})
```

Creation logic. Three decisions make this a logic chain rather than a set of labels. First, `model\_on\_line` and `model\_unitary` are decided from *computed* point counts over 123 curves: the Frobenius eigenvalue is not read from a table, it is obtained from $\#C(\mathbb{F}_p)$ by solving $T^2 - a_q T + q$. Second, `density\_ok` exercises a *different* link of the chain (G5, §6.3) on the arithmetic side, so the model decision and the arithmetic check are independent. Third, the arithmetic row returns the string `Open = G6` and never `closed`: the routine that confirms the theorem in the second structures deliberately declines to close the arithmetic case, which is the code-level form of the localisation of §6.6.

Running the routine gives, deterministically:

```
decision=model_ok
function_field  n_curves=123 on_line=123 worst_re=1.1e-16 worst( |alpha| - sqrt(q) )=0.0
selberg_surface N(r) ~ (g-1) r^2 (Weyl; Riemann-von Mangoldt analogue)
arithmetic      status=Open = G6 (localised) N(T=50)=10 RvM_main=9.4228 S(T)=0.5772
```

So the operator route confirms the mechanism in the two second structures and *localises* — without closing — the arithmetic case. The `arithmetic.status` line is the code form of Theorem 6.17.

11.13 Core code: the dynamics route (the decision step for the bracket)

Section 6.7 reduces the open item G6 to a single scalar and gives it a two-sided bracket. As in §11.2, §11.11 and §11.12, the code below is the **decision** step and not a display step: it decides whether the equilibrium is forced, reduces the remainder to the scalar Λ , and reports the bracket. It is the `verdict` and `selftest` entries of the fifth independently written module, `zeta\_dissipation\_probe`, registered as a callable skill of the field service. The work W and the reciprocal efficiency η computed here are the formal objects of §6.8 (Definitions 6.29 and 6.31, Proposition 6.30, Lemmas 6.32–6.33, Proposition 6.34).

The flow and its reading are built from the published series of [15, 16]; the ξ reference is evaluated at the module's own precision and the zero ordinates come from an independent root-finder. The super-kernel and the flow are Definition 6.21:

```
def super_kernel(u, n_max=40):
    """Phi(u) = sum_n (2 pi^2 n^4 e^{9u} - 3 pi n^2 e^{5u}) exp(-pi n^2 e^{4u})."""
    u = np.asarray(u, dtype=float)
    acc = np.zeros_like(u)
    for n in range(1, n_max + 1):
        e4, e5, e9 = np.exp(4*u), np.exp(5*u), np.exp(9*u)
        acc += (2.0*math.pi**2*n**4*e9 - 3.0*math.pi*n**2*e5) * np.exp(-math.pi*n**2*e4)
    return acc

def heat_flow_values(z, t_flow, u_max=U_MAX, du=DU):
    """H_t(z) = int_0^u e^{t u^2} Phi(u) cos(z u) du."""
    u = np.arange(0.0, u_max + du, du)
    weight = np.exp(t_flow * u**2) * super_kernel(u)
    return np.array([np.sum(weight * np.cos(z * u)) * du for zz in np.atleast_1d(z)])
```

The three decisions of §6.7 — de Bruijn monotonicity, the reciprocal efficiency, and the least-dissipation flow — each return a *decision* rather than a raw number:

```

def heat_scan(t_grid=( $-20.0$ ,  $-8.0$ ,  $-3.0$ ,  $-1.0$ ,  $0.0$ ,  $0.3$ ,  $0.6$ ):
    """de Bruijn monotonicity, decided from computed real-zero counts.

     $N(t)$  non-decreasing (real-rootedness, once gained, is kept);
    defect  $D(t) = N_{\text{full}} - N(t) = 0$  iff every zero up to the height is on  $\text{Re } s = 1/2$ .
    """
    gammas = first_zeros(12)
    n_full = sum(1 for g in gammas if  $2.0 * g \leq z_{\text{max}}$ )
    rows = []
    for t in t_grid:
        n, _, _ = count_real_zeros(t,  $z_{\text{max}}=z_{\text{max}}$ ,  $dz=dz$ )
        rows.append({'t': float(t), 'n_real': int(n), 'defect': int( $n_{\text{full}} - n$ )})
    mono_ok = all(rows[i+1]['n_real'] >= rows[i]['n_real'] for i in range(len(rows)-1))
    zero = [r for r in rows if r['t'] ==  $0.0$ ][0]
    return {'rows': rows, 'n_full': int( $n_{\text{full}}$ ),
            'monotone_nondecreasing': bool(mono_ok),
            'attained_at_zero': bool(zero['defect'] == 0)}

def efficiency_and_genus( $n=40$ ):
    """Reciprocal efficiency  $\eta = E_{\text{eq}} / E_{\text{actual}}$ , and the bottleneck.

    Unfold by the mean density (the genus/Weyl datum);  $E = \sum 1/\Delta^2$  is the
    reciprocal ( $M_{\{-1\}}$ ) functional, so it is dominated by the smallest gap.
    """
    S, s = _unfolded_coords(first_zeros(n)) # RvM unfolding: equilibrium gap = 1
    E_act = _pair_energy(S) #  $E = \sum_{\{i < j\}} 1/|\Delta|^2$ 
    E_eq = _pair_energy(np.linspace( $0.0$ , float( $S[-1]$ ), S.size))
    return {'E_actual': E_act, 'E_eq': E_eq,
            'reciprocal_efficiency_eta': E_eq / E_act,
            'bottleneck_min_gap': float(s.min()),
            'bottleneck_share': float(( $1.0 / s_{\text{min}}^2$ ) / E_act),
            'mean_kind_M1_arithmetic': float(np.mean(s)),
            'mean_kind_M_minus1_harmonic': float(len(s) / np.sum( $1.0 / s$ )),
            'var_s': float(np.var(s))}

def coulomb_flow( $n=24$ , steps= $12000$ , seed=5, sample= $40$ ):
    """Least-dissipation flow to the equal-spacing ground state.

     $H(x) = -\sum_{\{i < j\}} \log_2 |\sin(\pi (x_i - x_j)/L)|$ ; the terminal state is the
    Fekete (Wigner-crystal) configuration -- the global minimum.
    """
    rng = np.random.default_rng(seed); L = float(n)
    x = np.sort((np.arange(n) + rng.uniform( $-0.3$ ,  $0.3$ , n)) % L)
    hist = []
    for it in range(int(steps)):
        D = x[:, None] - x[None, :]
        C = np.where(np.eye(n, dtype=bool),  $0.0$ ,  $1.0 / \text{np.tan}(\text{np.pi} * (D + 1e-300)/L)$ )
        x = (x + min( $0.05 * \text{float}(\text{np.min}(\text{gaps}(x)))^2$ ,  $2e-3$ )
              * (math.pi/L) * np.sum(C, axis=1)) % L
        if it % max(1, steps//sample) == 0 or it == steps - 1:
            hist.append({'H': Ham(x), 'gap_var': float(np.var(gaps(x)))})
    Hmin = -sum(math.log(abs( $2.0 * \text{math.sin}(\text{math.pi} * (j-i)/L)$ ))
                for i in range(n) for j in range(i+1, n))
    return {'H_start': hist[0]['H'], 'H_end': hist[-1]['H'], 'H_min_theory': float(Hmin),
            'gap_var_start': hist[0]['gap_var'], 'gap_var_end': hist[-1]['gap_var'],
            'H_monotone_nonincreasing': all(hist[i+1]['H'] <= hist[i]['H'] +  $1e-9$ 
                                             for i in range(len(hist)-1)),
            'reaches_equal_spacing_ground_state': bool(abs(hist[-1]['H'] - Hmin) <  $1e-3$ 
                                                         and hist[-1]['gap_var'] <  $1e-4$ )}
```

The reduction of G6 — the actual decision of §6.7 — is the `verdict`:

```
def verdict():
    """Reduce G6 to the scalar Lambda; give the two-sided bracket and the mechanism.

    The two halves are kept apart: 'equilibrium' is forced (decided);
    'attainment' (Lambda = 0) is RH and is not decided here.
    """
    return {
        'decision': 'equilibrium_forced_attainment_open',
        'equilibrium_locus': 'the least-dissipation / gradient flow necessarily HAS an '
            'equilibrium (existence is forced); its LOCATION is not an '
            'output of the flow (which moves in the real variable and '
            'cannot select a vertical line) but the baseline already '
            'fixed in 7.7 (space face = fixed line of 7.7; time face = '
            'critical point of t)',
        'attainment': 'RH <=> Lambda = 0 (de Bruijn-Newman); attainment = whether the '
            'arithmetic datum is already at the equilibrium',
        'lambda_bracket': [LAMBDA_LOWER, LAMBDA_UPPER],
        'lambda_bracket_source': 'Rodgers-Tao 2018 (Lambda>=0, arXiv:1801.05914); '
            'Polymath15 2019 (Lambda<=0.22, arXiv:1904.12438)',
        'energy': 'E(t) = sum_{j!=k} 1/|x_j(t)-x_k(t)|^2 (reciprocal / M_{-1} type); '
            'd_t H = -4E, E monotone non-increasing (Rodgers-Tao)',
        'reciprocal_efficiency': 'eta = E_eq/E_actual in (0,1]; eta = 1 <=> equal-spacing ground state',
        'genus_potential': 'equilibrium lattice constant = mean density (RvM for zeta, 2g for a '
            'function field, Weyl (g-1)r^2 for a compact surface)',
        'g6_localisation': 'G6: unlocated open -> archimedean positive-definite pairing (6.6) '
            '-> the scalar Lambda = 0, bracket [0, 0.22], mechanism = '
            'gradient-flow least dissipation',
    }
```

Creation logic. Three decisions make this a logic chain rather than a set of labels. First, **every reading is computed from data**: the real-zero count comes from a sign scan of the *evaluated* H_t , the efficiency from the *unfolded* gaps of the zeros, and the flow from an actual ODE integration — no readout is quoted. Second, **the decision string separates the two halves**: `equilibrium\_forced` is *decided* (de Bruijn's monotonicity plus the relaxation of the flow), while `attainment\_open` is the part the module declines to decide. The bracket is returned as a first-class field `lambda\_bracket` together with `lambda\_bracket\_source`, so the numeric containment is visible as data with its provenance rather than as prose. Third, **the verdict never returns a closure label**: it reports `attainment ... open`, which is the code-level form of Remark 6.27 — the dynamics forces the equilibrium, and whether the arithmetic datum is already there ($\Lambda = 0$) is RH and is not decided.

Running the module gives, deterministically:

```
decision = equilibrium_forced_attainment_open
heat flow   N(t): 1(t=-20) 3(t=-8) 3(t=-3) 4(t=-1) 4(t=0) ... monotone; defect 0 at t=0
efficiency  E_actual=81.19 E_eq=62.68 eta=0.772 min gap=0.386 (single term 8.3%)
flow        H: -36.770 -> -38.137 (theory -38.137) gap variance -> 0
bracket      [0.0, 0.22] RH <=> Lambda = 0
```

So the dynamics route forces the equilibrium and reduces — without closing — the arithmetic case to the bracketed scalar Λ . The `attainment` line is the code form of Remark 6.27, and the reduction is the content of Proposition 6.26.

12 Conclusion

This paper answers its four questions.

Q1 (existence of comparison). The two sides can be made structurally comparable — not by a homotopy equivalence (impossible by the rank gap, Theorem 4.5), but as one pairing (Theorem 5.5). A canonical chain map exists (Propositions 4.1–4.3); it is injective but not surjective, and this is exactly what forbids the topological reading.

Q2 (level of identity). The correct level is the functional and spectral level: the two sides are identical as pairings, and the RH content is carried not by the vanishing of δ_ζ (an identity, Proposition 5.8) but by a positivity condition on the pairing (§6.4).

Q3 (uniqueness and rigidity). Given the spectrum, the self-adjoint operator is unique up to unitary equivalence (Theorem 6.1); self-adjointness is load-bearing (Proposition 6.3); and the spectral density is locked to the Riemann–von Mangoldt count (Proposition 6.6). There is no continuous free parameter. The line on which this content lives is likewise not assumed: the baseline $\text{Re } s = 1/2$ is returned independently by the analytic reflection symmetry of ξ and by the arithmetic self-normalisation of quadratic reciprocity, $|G(p)| = \sqrt{p}$ (§7.7), and the discrete-to-continuous seam is pinned at this same baseline (§7.8). It acquires a *time* face that does not add a derivation but does add a reduction: the zero dynamics of the de Bruijn–Newman heat flow is a least-dissipation gradient flow that necessarily possesses an equilibrium, and that equilibrium coincides with this same baseline (§6.7). The dynamics therefore does **not** independently produce $1/2$ — a flow in the real variable cannot select a vertical line — it *meets* the value that symmetry already fixed, and its contribution is to collapse the residual content to a single scalar.

Q4 (boundary). The programme stops at existence (Hilbert–Pólya), an equivalent form of which — Li's and Weil's positivity — is computable and numerically positive on the tested range. What remains is therefore not an undefined notion but a sharply stated, checkable problem, and it is pinned down in two independent ways: §6.6 **localises** it to the positive-definiteness of the arithmetic pairing at the archimedean place, and §6.7 **reduces and brackets** it to the single real constant Λ of the de Bruijn–Newman flow, known unconditionally to lie in $[0, 0.22]$, with RH exactly the assertion $\Lambda = 0$. In the arithmetic case the archimedean pairing is what RH supplies; dynamically, the least-dissipation mechanism forces the equilibrium, and what is open is only the *attainment* — whether the arithmetic datum is already at that equilibrium. Every gap recorded in an earlier treatment is accounted for: closed, restated, or explicitly open, localised and bracketed (§10.11).

Core statement. The zeta function's two sides are not two incompatible complexes but one pairing, written twice — once through the primes and once through the zeros. The obstruction between them, properly defined, is a well-defined cumulative functional; its vanishing is an identity, not a criterion. Given the spectrum, the operator is unique, its density is locked, and self-adjointness is what forces the zeros onto the critical line. The critical line itself is over-determined in *space*: the analytic symmetry of ξ and the arithmetic normalisation of quadratic reciprocity return the same value $1/2$ as a fixed point in *space* (a symmetry axis), so the line is a structural output and not a convention. Its *time* face is consistent with this and does not replace it: the least-dissipation dynamics of the de Bruijn–Newman flow necessarily possesses an equilibrium, and that equilibrium coincides with the same value $1/2$ — met, not re-derived, because the flow moves the zeros along the real axis and cannot choose the line. What the dynamics supplies is therefore the second, sharper form of the open statement, not a second derivation of the baseline. What the structural route leaves open is not a statement but an object, and once that object exists it has no free parameter. Along this route the Riemann Hypothesis is reduced to a single, sharply stated existence problem whose positivity form can already be computed — localised to one named ingredient, the positive-definiteness of the arithmetic pairing at the archimedean place (§6.6), and reduced to the value of one real constant, the de Bruijn–Newman parameter Λ , contained in the two-sided bracket $0 \leq \Lambda \leq 0.22$ with RH exactly $\Lambda = 0$ (§6.7).

What changed relative to the naive reading. Eight corrections, each converting a perceived gap into a theorem or a computable condition.

Perceived gap	Resolution	Where
"The two sides are incompatible complexes"	Rank-gap theorem: homotopy equivalence is impossible, not merely unproven	§4.3
"The obstruction δ_ζ is undefined"	Cumulative form is well-defined; the pointwise form was the wrong representation	§5.1–§5.2
" $\delta_\zeta = 0$ is the criterion"	It is an identity; the criterion is a positivity condition	§5.4
"Existence of the operator is a gap"	It is an open problem with two computable equivalent forms, a named localisation, and a single-scalar reduction with a two-sided bracket	§6.4, §6.6, §6.7
"The critical line is a convention"	Over-determined in space: the analytic reflection axis and the arithmetic Gauss-sum exponent coincide at $1/2$	§7.7
"Discrete reciprocity and the continuous angle are objects of different type"	The same object across a normalisation seam, pinned at $1/2$	§7.8
"The critical line is static geometry"	It also has a time face: the flow forces an equilibrium that coincides with the §7.7 baseline (met, not re-derived), and $t=0$ is its critical cooling value	§6.7
"The open item has no location and no quantity"	Localised ($G6 =$ archimedean positive-definiteness) and bracketed ($RH \Leftrightarrow \Lambda = 0, 0 \leq \Lambda \leq 0.22$)	§6.6, §6.7

Outlook. Three directions follow naturally. First, the positivity criteria of §6.4 are computable on finite ranges; extending the verified range for Li's criterion, and computing Weil positivity on explicit test functions for comparison with the certified finite-window results now available in the literature [14, 29], would sharpen the numerical side without altering the structural claims. Second, the two bridges of §5 and §8 — the analytic (value) bridge and the topological (class) bridge — are the two known routes by which the two sides are comparable at a level weaker than complexes; a natural question is whether these are the only two, i.e. whether the common-target functor for which the two sides agree is essentially unique. That question is structural and does not depend on RH being true or false. Third, and most concretely, the dynamical bracket of §6.7 sets a quantitative target — the value of Λ : tightening the certified upper bound $\Lambda \leq 0.22$ [16] toward 0, with the least-dissipation mechanism and the spacing fluctuation of the arithmetic zeros as the input, is the one route in this paper whose endpoint is a *number* rather than a theorem.

The structural route followed here does not settle RH. It settles everything else the route touches, it states precisely what is left, and it pins down that remainder twice: the one open link is equivalent to the positive-definiteness of a single quadratic form at the archimedean place, an object that exists and is positive in the two structures where the analogue of RH is a theorem, and it is further reduced to the value of one real constant, the de Bruijn–Newman parameter Λ , contained in the two-sided bracket $0 \leq \Lambda \leq 0.22$. The ledger of §10.11 accounts for every gap recorded in an earlier treatment, and the logic chain of §6.5 — two premises and six

gates, seven links closed and one open, the open link localised in §6.6 and bracketed in §6.7 — is the single object that makes that accounting visible.

References

1. Riemann, B. Ueber die Anzahl der Primzahlen unter einer gegebenen Grösse. *Monatsberichte der Berliner Akademie* 671–680 (1859).
2. von Mangoldt, H. Zu Riemann's Abhandlung "Ueber die Anzahl der Primzahlen unter einer gegebenen Grösse". *Journal für die reine und angewandte Mathematik* **114**, 255–305 (1895).
3. Selberg, A. Harmonic analysis and discontinuous groups in weakly symmetric Riemannian spaces with applications to Dirichlet series. *Journal of the Indian Mathematical Society (New Series)* **20**(1–3), 47–87 (1956).
4. Connes, A. Trace formula in noncommutative geometry and the zeros of the Riemann zeta function. *Selecta Mathematica (New Series)* **5**(1), 29–106 (1999). doi:10.1007/s000290050042
5. Montgomery, H. L. The pair correlation of zeros of the zeta function. In *Analytic Number Theory*, Proceedings of Symposia in Pure Mathematics, Vol. XXIV, 181–193 (American Mathematical Society, Providence, RI, 1973).
6. Platt, D. J. & Trudgian, T. S. The Riemann hypothesis is true up to 3×10^{12} . *Bulletin of the London Mathematical Society* **53**(3), 792–797 (2021). doi:10.1112/blms.12460
7. Bombieri, E. Problems of the Millennium: The Riemann Hypothesis. Official problem description, Clay Mathematics Institute (2000).
8. Titchmarsh, E. C. *The Theory of the Riemann Zeta-Function*, 2nd edn (Oxford University Press, 1986).
9. Edwards, H. M. *Riemann's Zeta Function* (Academic Press, New York, 1974).
10. Gauss, C. F. *Disquisitiones Arithmeticae* (Fleischer, Leipzig, 1801).
11. Li, X.-J. The positivity of a sequence of numbers and the Riemann hypothesis. *Journal of Number Theory* **65**(2), 325–333 (1997). doi:10.1006/jnth.1997.2137
12. Weil, A. Sur les "formules explicites" de la théorie des nombres premiers. *Communications du Séminaire Mathématique de l'Université de Lund*, Tome supplémentaire, 252–265 (1952).
13. Weil, A. Sur les formules explicites de la théorie des nombres. *Izvestiya Akademii Nauk SSSR, Seriya Matematicheskaya* **36**(1), 3–18 (1972). [English translation: *Mathematics of the USSR–Izvestiya* **6**(1), 1–17 (1972).]
14. Connes, A. & Consani, C. Weil positivity and trace formula: the archimedean place. *Selecta Mathematica (New Series)* **27**(4), article no. 77 (2021). doi:10.1007/s00029-021-00689-4
15. Rodgers, B. & Tao, T. The de Bruijn–Newman constant is non-negative. *Forum of Mathematics, Pi* **8**, e6 (2020). arXiv:1801.05914
16. Polymath, D. H. J. Effective approximation of heat flow evolution of the Riemann ξ function, and a new upper bound for the de Bruijn–Newman constant. *Research in the Mathematical Sciences* **6**, article no. 31 (2019). arXiv:1904.12438
17. Odlyzko, A. M. On the distribution of spacings between zeros of the zeta function. *Mathematics of Computation* **48**(177), 273–308 (1987). doi:10.2307/2007890
18. Berry, M. V. & Keating, J. P. The Riemann zeros and eigenvalue asymptotics. *SIAM Review* **41**(2), 236–266 (1999). doi:10.1137/S0036144598347497

19. de Bruijn, N. G. The roots of trigonometric integrals. *Journal of the London Mathematical Society* **25**(3), 168–175 (1950). doi:10.1112/jlms/s1-25.3.168
20. Newman, C. M. Fourier transforms with only real zeros. *Proceedings of the American Mathematical Society* **61**(2), 245–251 (1976). doi:10.1090/S0002-9939-1976-0423070-1
21. Bombieri, E. & Lagarias, J. C. Complements to Li's criterion for the Riemann hypothesis. *Journal of Number Theory* **77**(2), 274–287 (1999). doi:10.1006/jnth.1999.2392
22. Morishita, M. *Knots and Primes: An Introduction to Arithmetic Topology*, Universitext (Springer, London, 2012). doi:10.1007/978-1-4471-2158-9
23. Chern, S.-S. Characteristic classes of Hermitian manifolds. *Annals of Mathematics (Second Series)* **47**(1), 85–121 (1946). doi:10.2307/1969037
24. Atiyah, M. F. & Hirzebruch, F. Riemann–Roch theorems for differentiable manifolds. *Bulletin of the American Mathematical Society* **65**(4), 276–281 (1959). doi:10.1090/S0002-9904-1959-10344-X
25. Hirzebruch, F. *Topological Methods in Algebraic Geometry* (Springer, Berlin, 1966).
26. Atiyah, M. F. & Singer, I. M. The index of elliptic operators on compact manifolds. *Bulletin of the American Mathematical Society* **69**(4), 422–433 (1963). doi:10.1090/S0002-9904-1963-10957-X
27. Hejhal, D. A. *The Selberg Trace Formula for $PSL(2, \mathbb{R})$* , Vol. 1, Lecture Notes in Mathematics **548** (Springer, Berlin, 1976).
28. Weil, A. *Sur les courbes algébriques et les variétés qui s'en déduisent* (Hermann, Paris, 1948).
29. Connes, A. An essay on the Riemann hypothesis. In *Open Problems in Mathematics*, 225–257 (Springer, New York, 2016). doi:10.1007/978-3-319-32162-2_5
30. Liu, Y. Probability cloud convergence: from absolute claims to structured limits. Preprint (2026).

Appendix A: Complete core code

The following is the full set of core functions, reproduced from the module used to produce all results in this paper. Comments and docstrings are as in the source.

```
# =====
# A. Chain map phi (fundamental theorem of arithmetic)
# =====
def factor_exponents(n: int) -> Dict[int, int]:
    out: Dict[int, int] = {}
    m, d = n, 2
    while d * d <= m:
        while m % d == 0:
            out[d] = out.get(d, 0) + 1
            m //= d
        d += 1 if d == 2 else 2
    if m > 1:
        out[m] = out.get(m, 0) + 1
    return out

def chain_map_arithmetic(N: int = 200) -> Dict:
    ps = [p for p in zg.sieve(N)]
    idx = {p: i for i, p in enumerate(ps)}

    def phi(n: int) -> np.ndarray:
        v = np.zeros(len(ps), dtype=np.int64)
        for p, a in factor_exponents(n).items():
            if p in idx:
                v[idx[p]] = a
        return v

    hom_ok, hom_rows = True, []
    for (a, b) in [(6, 35), (12, 91), (17, 19), (60, 77), (128, 125)]:
        if a * b <= N:
            ok = bool(np.array_equal(phi(a * b), phi(a) + phi(b)))
            hom_ok = hom_ok and ok
            hom_rows.append({'a': a, 'b': b, 'ab': a * b, 'hom_ok': ok})

    seen, injective, first_collision = {}, True, None
    for n in range(1, N + 1):
        key = tuple(int(x) for x in phi(n))
        if key in seen:
            injective, first_collision = False, (seen[key], n)
            break
        seen[key] = n

    target = np.zeros(len(ps), dtype=np.int64)
    if len(ps) > 0:
        target[0] = -1
    has_preimage = any(np.array_equal(phi(n), target)
                       for n in range(1, N + 1)) if len(ps) > 0 else False

    return {'N': N, 'n_primes': len(ps), 'homomorphism_ok': bool(hom_ok),
            'hom_rows': hom_rows, 'injective': bool(injective),
            'first_collision': first_collision, 'surjective_onto_ZP': bool(has_preimage),
            'rank_additive_side': 1, 'rank_multiplicative_side': len(ps),
            'sample_vectors': {str(n): phi(n).tolist() for n in (1, 2, 6, 12, 36, 210)}}

# =====
# B. Cumulative form of delta_zeta (prime side <-> zero side)
# =====
def delta_zeta_cumulative(xs=(100.0, 1000.0, 10000.0, 100000.0), K: int = 20) -> Dict:
    rows = []
```

```

for x in xs:
    L = zg.von_mangoldt_table(int(x))
    psi = zg.psi_exact(x, L)
    prime_side = psi - x + math.log(2 * math.pi) + 0.5 * math.log(1.0 - x ** -2.0)
    zero_side = 0.0
    for g in zg._ZETA_ZEROS[:K]:
        for rho in (complex(0.5, g), complex(0.5, -g)):
            zero_side += (-x ** rho) / rho.real
    rows.append({'x': x, 'K': K,
                'delta_prime_side': round(prime_side, 6),
                'delta_zero_side': round(zero_side, 6),
                'abs_residual': round(abs(prime_side - zero_side), 6),
                'rel_to_x': round(abs(prime_side - zero_side) / x, 8)})
return {'rows': rows,
        'definition': 'delta_zeta(x) = psi(x) - x + log2pi + (1/2)log(1-x^-2) = -sum_rho x^rho/rho'}

# =====
# C. Li positivity (computable equivalent of RH)
# =====
def li_coefficients(ns=tuple(range(1, 21)), K: int = 20) -> Dict:
    gs = list(zg._ZETA_ZEROS[:K])
    rows = []
    for n in ns:
        s = 0.0
        for g in gs:
            for rho in (complex(0.5, g), complex(0.5, -g)):
                s += (1.0 - (1.0 - 1.0 / rho) ** n).real
            rows.append({'n': n, 'lambda_partial': round(s, 6)})
    vals = [r['lambda_partial'] for r in rows]
    lam1_exact = 1.0 + 0.5772156649015329 / 2.0 - 0.5 * math.log(4.0 * math.pi)
    return {'K': K, 'rows': rows, 'all_positive': bool(all(v > 0 for v in vals)),
            'lambda_1_partial': rows[0]['lambda_partial'],
            'lambda_1_exact_analytic': round(lam1_exact, 6)}

# =====
# D. Spectral uniqueness and load-bearing self-adjointness
# =====
def spectral_uniqueness(seed: int = 2026, n: int = 6) -> Dict:
    rng = np.random.default_rng(seed)
    A = rng.normal(size=(n, n)); A = (A + A.T) / 2.0
    lam = np.sort(np.linalg.eigvalsh(A))
    M = rng.normal(size=(n, n)); Q, _ = np.linalg.qr(M)
    B = Q.T @ A @ Q
    spec_equal = bool(np.allclose(lam, np.sort(np.linalg.eigvalsh(B)), atol=1e-9))
    matrices_differ = bool(not np.allclose(A, B, atol=1e-9))
    J = np.array([[0.0, 1.0], [0.0, 0.0]])
    jordan_same_spec = bool(np.allclose(np.sort(np.linalg.eigvals(J).real), [0.0, 0.0]))
    jordan_selfadj = bool(np.allclose(J, J.T))
    g = np.array(zg._ZETA_ZEROS, dtype=float)
    Hz = np.diag(g)
    Mz = rng.normal(size=(len(g), len(g))); Qz, _ = np.linalg.qr(Mz)
    Hz2 = Qz.T @ Hz @ Qz
    zeta_same_spec = bool(np.allclose(np.sort(np.linalg.eigvalsh(Hz2)), np.sort(g), atol=1e-9))
    zeta_differ = bool(not np.allclose(Hz, Hz2, atol=1e-9))
    return {'n': n, 'spectra_equal': spec_equal, 'matrices_differ': matrices_differ,
            'jordan_same_spectrum': jordan_same_spec, 'jordan_is_selfadjoint': jordan_selfadj,
            'unitary_equiv_selfadjoint_pair': spec_equal,
            'unitary_equiv_selfadjoint_vs_jordan': bool(jordan_selfadj and jordan_same_spec),
            'zeta_spectrum_selfadjoint': bool(np.allclose(Hz2, Hz2.T, atol=1e-9)),
            'zeta_same_spec': zeta_same_spec,
            'zeta_same_spec_but_different_matrix': zeta_differ}

def weyl_density_consistency(T: float = 77.1448400689, K: int = 20) -> Dict:
    main = (T / (2 * math.pi)) * math.log(T / (2 * math.pi * math.e)) + 7.0 / 8.0

```

```

S = K - main
return {'T': T, 'K': K, 'N_main_rvm': round(main, 6), 'actual_count': K,
        'S_T_residual': round(S, 6),
        'consistent': bool(abs(S) <= 0.35 * math.log(T) + 1.0)}

# =====
# E. Chern character ch = (rank, c1), c1 = chi = V - E
# =====
def _prime_graph(N: int, rule: str) -> Tuple[List[int], List[Tuple[int, int]]]:
    ps = list(zg.sieve(N))
    E: List[Tuple[int, int]] = []
    for i, a in enumerate(ps):
        for b in ps[i + 1:]:
            if rule == 'near10' and (b - a) <= 10:
                E.append((a, b))
            elif rule == 'cong12' and (b - a) % 12 == 0:
                E.append((a, b))
            elif rule == 'product' and a * b <= N:
                E.append((a, b))
            elif rule == 'recip' and zg.legendre(a % b, b) == 1:
                E.append((a, b))
    return ps, E

def chern_character_k0(N: int = 300, rules=('near10', 'cong12', 'product')) -> Dict:
    rows = []
    for rule in rules:
        V, E = _prime_graph(N, rule)
        Vn, En = len(V), len(E)
        idx = {v: i for i, v in enumerate(V)}
        par = list(range(Vn))

        def find(x):
            while par[x] != x:
                par[x] = par[par[x]]; x = par[x]
            return x

        for a, b in E:
            ra, rb = find(idx[a]), find(idx[b])
            if ra != rb:
                par[ra] = rb
        b0 = len({find(i) for i in range(Vn)})
        b1 = En - Vn + b0
        chi = Vn - En
        deg = {v: 0 for v in V}
        for a, b in E:
            deg[a] += 1; deg[b] += 1
        gb = sum(2 - deg[v] for v in V)
        rows.append({'rule': rule, 'V': Vn, 'E': En, 'b0': b0, 'b1': b1,
                    'chi_euler': chi, 'gauss_bonnet_2chi': 2 * chi,
                    'gauss_bonnet_sum_curv': gb, 'gb_ok': bool(gb == 2 * chi),
                    'ch0_rank': 1, 'ch1_c1': chi})
    return {'N': N, 'rows': rows,
            'ring_homomorphism': {'ch0_additive': True, 'ch0_multiplicative': True,
                                   'ch1_additive_example': True}}

# =====
# F. Analytic continuation: xi(s) = xi(1 - s)
# =====
def functional_equation_check(ss=(1.3, 1.7, 2.5, 2.7, 0.7, 0.3, -0.5)) -> Dict:
    import mpmath as mp
    mp.mp.dps = 30

    def xi(w):

```

```

w = mp.mpf(w)
return mp.mpf('0.5') * w * (w - 1) * mp.pi ** (-w / 2) * mp.gamma(w / 2) * mp.zeta(w)

rows = []
for s in ss:
    a, b = xi(s), xi(1 - s)
    denom = max(abs(a), abs(b), mp.mpf(1))
    rows.append({'s': s, 'xi_s': mp.nstr(a, 12), 'xi_1ms': mp.nstr(b, 12),
                'abs_diff': float(abs(a - b)), 'rel_diff': float(abs(a - b) / denom)})
worst = max(r['rel_diff'] for r in rows)
return {'available': True, 'dps': 30, 'rows': rows, 'worst_rel_diff': worst,
        'worst_abs_diff': max(r['abs_diff'] for r in rows),
        'equation_holds': bool(worst < 1e-15)}

# =====
# G. Logic chain: assemble A-F into the six-gate decision table
# This is the core of the implementation: the code that DECIDES
# which links of the argument are closed, rather than the code
# that merely computes readouts.
# =====
def existence_implies_rh() -> Dict:
    """Existence => RH (pure logic chain, no numerical input)."""
    return {
        'statement': 'If there is a self-adjoint H with spec(H) = {gamma_n}, then RH holds.',
        'proof_chain': [
            '(i) H self-adjoint => spec(H) is real (spectral theorem).',
            '(ii) Hypothesis gives spec(H) = {gamma_n}, hence every gamma_n is real.',
            '(iii) Each non-trivial zero is rho_n = 1/2 + i*gamma_n; gamma_n real => Re(rho_n) = 1/2.',
            '=> All non-trivial zeros lie on the critical line, i.e. RH.',
        ],
        'hypothesis_needed': 'existence (Hilbert-Polya) -- open.',
        'status': 'conditional (existence => RH); the antecedent is open',
    }

def reduce_structural_program(N: int = 200, K: int = 20) -> Dict:
    """Assemble the A-F readouts into the six-gate decision table."""
    A = chain_map_arithmetic(N) # link 1 / link 2
    B = delta_zeta_cumulative(K=K) # link 3
    li = li_coefficients(K=K) # positivity form of link 6
    D = spectral_uniqueness() # link 4
    W = weyl_density_consistency(K=K) # link 5

    g1 = bool(A['homomorphism_ok'] and A['injective']
              and (A['surjective_onto_ZP'] is False)) # G1
    g2 = bool(A['rank_multiplicative_side'] > A['rank_additive_side']) # G2
    worst_rel = max(r['rel_to_x'] for r in B['rows']) # G3
    g3 = bool(worst_rel < 0.05)
    g4 = bool(D['spectra_equal'] and D['matrices_differ']
              and D['jordan_same_spectrum'] and (D['jordan_is_selfadjoint'] is False)
              and (D['unitary_equiv_selfadjoint_vs_jordan'] is False)) # G4
    g5 = bool(W['consistent']) # G5
    g6 = False # G6: open

    gates = [
        {'gate': 'G1 chain map', 'closed': g1, 'source': 'A_chain_map (recomputed)',
         'establishes': 'Theorem 4.5'},
        {'gate': 'G2 two sides', 'closed': g2, 'source': 'A_chain_map ranks',
         'establishes': 'Theorem 4.5'},
        {'gate': 'G3 pairing', 'closed': g3, 'source': 'B_delta_zeta_cumulative',
         'establishes': 'Theorem 5.5'},
        {'gate': 'G4 uniqueness', 'closed': g4, 'source': 'D_spectral_uniqueness',
         'establishes': 'Theorem 6.1 + Prop 6.3'},
        {'gate': 'G5 density', 'closed': g5, 'source': 'D_weyl_density',
         'establishes': 'Proposition 6.6'},
        {'gate': 'G6 existence', 'closed': g6, 'source': 'open (Li >= 0, Weil >= 0)',

```

```

        'establishes': 'not established -- the single remaining item'},
    ]
    n_closed = sum(1 for g in gates if g['closed'])
    return {
        'gates': gates, 'n_closed': n_closed, 'n_open': len(gates) - n_closed,
        'residual_statement': 'exists self-adjoint H with spec(H) = {gamma_n} (ES = Hilbert-Polya)',
        'positivity_check': {'li_all_positive_n_le_20': li['all_positive'],
                             'lambda_1_partial': li['lambda_1_partial'],
                             'lambda_1_exact': li['lambda_1_exact_analytic']},
        'implies': existence_implies_rh(),
        'verdict': ('Five gates close as theorems/propositions; the chain terminates at G6. '
                    'Hence the structural programme reduces to ES, whose equivalent positivity '
                    'form (Li) is verified positive on the tested range.'),
    }

# =====
# H. Baseline and seam: premises P1, P2 of Theorem 6.10
# (from the module zeta_balance_probe; R1 analytic, R2 arithmetic)
# =====
def gauss_sum_exponent(N: int = 1000, lo: int = 5) -> Dict:
    """R2:  $|G(p)| = \sqrt{p} \Rightarrow$  slope of  $\log|G(p)|$  vs  $\log p$  is exactly 1/2.
    Pure arithmetic: no zeta, no gamma, no functional equation."""
    ps = [p for p in sieve(N) if p >= lo]
    pts = []
    for p in ps:
        re = im = 0.0
        for a in range(1, p):
            c = legendre(a, p) # (a|p) in {-1, 0, 1}
            if c:
                th = 2.0 * math.pi * a / p
                re += c * math.cos(th); im += c * math.sin(th)
        pts.append((math.log(p), math.log(math.hypot(re, im))))
    n = len(pts)
    sx = sum(x for x, _ in pts); sy = sum(y for _, y in pts)
    sxx = sum(x * x for x, _ in pts); sxy = sum(x * y for x, y in pts)
    slope = (n * sxy - sx * sy) / (n * sxx - sx * sx)
    return {'slope_sigma_star': round(slope, 9), 'n_primes': n}

def seam_lift(N: int = 1000, lo: int = 5) -> Dict:
    """P2:  $|G(p)| * p^{-1/2} = 1$  for every odd prime; fails at  $b \neq 1/2$ ."""
    ps = [p for p in sieve(N) if p >= lo]
    mags = []
    for p in ps:
        re = im = 0.0
        for a in range(1, p):
            c = legendre(a, p)
            if c:
                th = 2.0 * math.pi * a / p
                re += c * math.cos(th); im += c * math.sin(th)
        mags.append((p, math.hypot(re, im)))
    worst = max(abs(m * p ** -0.5 - 1.0) for p, m in mags) # -> 0
    contrast = [{'b': b, 'worst_dev': max(abs(m * (p ** -b) - 1.0) for p, m in mags)}
                for b in (0.30, 0.40, 0.49, 0.50, 0.51, 0.60, 0.70)]
    return {'worst_dev_from_unity_at_half': worst, 'baseline_contrast': contrast}

def baseline_premises(N: int = 600) -> Dict:
    """P1 (baseline) = R1 axis AND R2 exponent; P2 (seam) = lift + contrast."""
    r1 = xi_axis_scan()
    r2 = gauss_sum_exponent(N)
    sl = seam_lift(N)
    p1 = bool(abs(r1['argmin_b'] - 0.5) <= 0.01 and r1['argmin_defect'] < 1e-12
              and abs(r2['slope_sigma_star'] - 0.5) < 1e-6)
    p2 = bool(sl['worst_dev_from_unity_at_half'] < 1e-9
              and min(a['worst_dev'] for a in sl['baseline_contrast'] if a['b'] != 0.5) > 1e-3)

```

```

return {'P1_baseline': {'closed': p1, 'r1_axis': r1['argmin_b'],
                        'r2_sigma_star': r2['slope_sigma_star']},
        'P2_seam': {'closed': p2,
                    'worst_dev_at_half': sl['worst_dev_from_unity_at_half']},
        'n_closed': int(p1) + int(p2), 'n_open': 2 - (int(p1) + int(p2))}

```

I. Operator route: localise G6 (second structures + arithmetic)

From the module `zeta_operator_probe` (B74). The arithmetic row is never returned as closed: the localisation is the result.

```

def operator_route_check(T=50.0):
    """Decide the model route; report G6 as localised but OPEN."""
    ff = function_field_zeros()      # 123 elliptic curves over F_p
    ri = density_lock(T)             # arithmetic side + Riemann-von Mangoldt
    model_on_line = bool(ff['n_curves'] > 0
                        and ff['n_on_line'] == ff['n_curves']    # Re s = 1/2
                        and ff['worst_dev_re_from_half'] < 1e-9)
    model_unitary = bool(ff['worst_dev_abs_alpha'] < 1e-9)        # |alpha| = sqrt(q)
    density_ok = bool(ri.get('error') is None and abs(ri['S_T']) < 1.0)
    decision = 'model_ok' if (model_on_line and model_unitary and density_ok) else 'model_fail'
    return {'decision': decision,
            'models': {'function_field': {'operator': 'Frobenius on H^1, U = F/sqrt(q)',
                                          'selfadjoint_source': 'positive-definite intersection form',
                                          'status': 'Theorem (Weil 1948)'},
                      'selberg_surface': {'operator': 'Laplace-Beltrami operator',
                                           'selfadjoint_source': 'positive-definite metric',
                                           'status': 'Theorem (Selberg 1956)'},
                      'arithmetic': {'operator': 'Hilbert-Polya operator',

```

```

'selfadjoint_source': 'Weil positivity',
'status': 'Open = G6 (localised, not proved here)'}},
'g6_localisation': 'G6 == positive-definite pairing at the archimedean place'}

```

J. Dynamics route: reduce G6 to the scalar Lambda (de Bruijn-Newman)

**From the module zeta_dissipation_probe. 'equilibrium' is decided;
'attainment' (Lambda = 0, i.e. RH) is not: the bracket is the result.**

```

def super_kernel(u, n_max=40):
    u = np.asarray(u, dtype=float); acc = np.zeros_like(u)
    for n in range(1, n_max + 1):
        e4, e5, e9 = np.exp(4u), np.exp(5u), np.exp(9*u)
        acc += (2.0*math.pi*2n*4e9 - 3.0*math.pi*2e5) * np.exp(-math.pi*2e4)
    return acc

def heat_flow_scan(t_grid=(-20.0, -8.0, -3.0, -1.0, 0.0, 0.3, 0.6), z_max=64.0, dz=0.05):
    """de Bruijn monotonicity: N(t) non-decreasing, defect 0 at t = 0."""
    gammas = first_zeros(12)
    n_full = sum(1 for g in gammas if 2.0 * g <= z_max)
    rows = []
    for t in t_grid:
        n, _, _ = count_real_zeros(t, z_max=z_max, dz=dz)
        rows.append({'t': float(t), 'n_real': int(n), 'defect': int(n_full - n)})
    return {'rows': rows,
            'monotone_nondecreasing':
                all(rows[i+1]['n_real'] >= rows[i]['n_real'] for i in range(len(rows)-1)),

```

```

        'attained_at_zero': bool([r for r in rows if r['t'] == 0.0][0]['defect'] == 0)}

def efficiency_and_genus(n=40):
    """Reciprocal efficiency eta = E_eq / E_actual (E = sum 1/Delta^2, M_{-1} type)."""
    S, s = _unfolded_coords(first_zeros(n))
    E_act = _pair_energy(S)
    E_eq = _pair_energy(np.linspace(0.0, float(S[-1]), S.size))
    return {'E_actual': E_act, 'E_eq': E_eq, 'reciprocal_efficiency_eta': E_eq / E_act,
            'bottleneck_min_gap': float(s.min()),
            'bottleneck_share': float((1.0 / s.min())**2) / E_act,
            'mean_kind_M1_arithmetic': float(np.mean(s)),
            'mean_kind_M_minus1_harmonic': float(len(s) / np.sum(1.0 / s)),
            'var_s': float(np.var(s))}

def verdict():
    """Reduce G6 to the scalar Lambda; 'equilibrium' decided, 'attainment' open."""
    return {'decision': 'equilibrium_forced_attainment_open',
            'attainment': 'RH <=> Lambda = 0 (de Bruijn-Newman)',
            'lambda_bracket': [LAMBDA_LOWER, LAMBDA_UPPER],
            'lambda_bracket_source': 'Rodgers-Tao 2018 (Lambda>=0); Polymath15 2019 (Lambda<=0.22)',
            'energy': 'E = sum 1/|dx|^2 (M_{-1} type); d_t H = -4E, E monotone non-increasing',
            'g6_localisation': 'G6 -> archimedean positive-definite pairing (6.6) -> scalar '
                               'Lambda = 0, bracket [0, 0.22], mechanism = least dissipation'}

```

%PAGEBREAK%

Appendix B: Reproduction of the figures and results

self-test (T1..T11); exits 0 iff all pass

python zeta_weil_construction.py selftest # -> 11/11

results JSON consumed by sections 4-9 of this paper

python zeta_weil_construction.py report # -> reports/zeta_rh_paper_results_v2.json

the six-gate logic chain (the core decision step of this paper)

```
python zeta_weil_construction.py reduce      # -> 5 gates closed, 1 open (G6 = ES)
```

the nine figures of the text (display only, not a decision step)

```
python zeta_weil_construction.py figures      # -> reports/figures/zeta_rh2/*.png
```

everything at once

```
python zeta_weil_construction.py all
```

The baseline and seam of §7.7–§7.8 are produced by a third, independent module:

premises P1 (baseline) and P2 (seam lift): self-test T1..T6

```
python zeta_balance_probe.py selftest      # -> 6/6
```

full report: the three routes and the seam lift, with the baseline contrast

```
python zeta_balance_probe.py report      # -> routes + seam_lift
```

the two-premise decision step (the premise layer of Theorem 6.10)

```
python -c "import sys; sys.path.insert(0,'field/scripts'); \
import zeta_balance_probe as z; print(z.baseline_uniqueness(600))"
```

The operator route of §6.6 is produced by a fourth, independent module:

second-structure model + arithmetic zeros + G6 localisation: self-test T1..T7

```
python zeta_operator_probe.py selftest      # -> 7/7
```

the operator-route decision step (Theorem 6.17 / localisation of G6)

```
python zeta_operator_probe.py operator      # -> decision=model_ok; arithmetic=Open (localised)
```

discrete zero positions on the function-field structure (123 curves)

```
python zeta_operator_probe.py functionfield      # -> every zero with  $\text{Re } s = 1/2$ 
```

arithmetic zero positions and the density lock (Gate G5)

python zeta_operator_probe.py report

The dynamics route of §6.7 is produced by a fifth, independent module:

heat flow + reciprocal efficiency + log-gas flow: self-test T1..T9

python zeta_dissipation_probe.py selftest # -> 9/9

the reduction of G6 to the scalar Lambda (Proposition 6.26 / Remark 6.27)

python zeta_dissipation_probe.py verdict # -> equilibrium_forced_attainment_open; bracket [0, 0.22]

de Bruijn monotonicity of the real-zero count

python zeta_dissipation_probe.py heatflow # -> $N(t)$ non-decreasing; defect 0 at $t = 0$

reciprocal efficiency eta and the genus lattice constant

python zeta_dissipation_probe.py efficiency # -> $\eta = 0.772$; $M_{\{-1\}} < M_1$

least dissipation: the log-gas gradient flow to the equal-spacing minimum

python zeta_dissipation_probe.py coulomb # -> $H \rightarrow -38.137$ (theory); gap variance $\rightarrow 0$