

Structural Arithmetic Geometry

Part III of *Arithmetic Positivity Program*

Matthew Long
The YonedaAI Collaboration
YonedaAI Research Collective, Chicago, Illinois
matthew@yonedaai.com

September 2026

Abstract

The analytic, spectral, cohomological, and operator formulations of the Riemann hypothesis suggest a common structural explanation, but equivalence of propositions does not provide maps between their mathematical objects. We organize a research program around explicit realizations of the classical Weil pairing. Relative to a fixed test space and Hermitian form, we describe the universal positive realization and prove its factorization property. Its existence is equivalent to positivity, so universality alone supplies no independent arithmetic input. We then distinguish this certificate from a geometric construction with a primitive intersection pairing, and from a spectral realization with an adjoint weight identity. Elementary counterexamples show why spectral purity without semisimplicity, positivity after a shift, and small eigenvector residuals without gap control cannot supply the missing bridges. For finite approximation programs, we specify the comparison maps, normalization, topology, and sign conditions required to pass to Weil positivity or to a holomorphic zero limit. The resulting obstruction is a concrete arithmetic comparison problem together with route-specific selection and convergence obligations. General transfer statements and several counterexamples have Lean implementations; the construction over the arithmetic base and the Riemann hypothesis remain open.

Subject classification: 11M26, 46C05, 47A75. *Keywords:* Weil positivity, arithmetic geometry, polarisation, spectral realization, formalization.

Contents

1	A program organized by maps	2
1.1	The fixed analytic target	3
1.2	Requirements for a structural explanation	3
2	Positive realizations and relative minimality	4
2.1	Weaker certificates and their limits	5
3	The geometric comparison as a pullback	5
3.1	Logical strength of the image-restricted inequality	6
3.2	Composition and descent	6
4	A spectral bridge with an explicit domain	7
4.1	From purity to an invariant metric	8

5	Polarisation from a positive Weil realization	8
5.1	The radical and the intended carrier	9
6	Limits of certificates on one test space	9
6.1	The quantifier order	10
6.2	Subtraction and normalization	10
7	The ground-state route and its smaller residual	11
7.1	An explicit sufficient norm estimate	11
7.2	Rayleigh bounds and ground-state selection	11
7.3	The holomorphic zero-transfer obligation	12
8	A network of bridges with typed obligations	12
8.1	A disciplined meaning of the common obstruction	13
9	Formal coverage and unproved arithmetic instances	13
9.1	Declarations versus theorem instances	13
9.2	The next independently testable questions	14
10	Finite Gram data and compatible limits	14
10.1	Consistent finite kernels	15
10.2	Direct limits of inner product spaces	15
10.3	Operators on the compatible limit	16
11	Arithmetic operator comparisons	16
12	Research outcome	17
A	Small countermodels	17
A.1	Dropping the comparison	17
A.2	Dropping injectivity	18
A.3	Dropping nonnegativity before quotienting	18
A.4	Dropping semisimplicity	18
A.5	Dropping continuity at a dense core	18
A.6	Dropping the gap scale	18
A.7	Dropping holomorphic control	18
B	An explicit comparison checklist	18
C	Logical strength of an unrestricted residual	19

1 A program organized by maps

The Arithmetic Positivity Program asks for structures that explain why several criteria for the Riemann hypothesis concern the same phenomenon. Its starting point is the classical Weil pairing on a specified space of test functions. Its desired output is an independently constructed arithmetic object equipped with maps that recover that pairing, its symmetries, and its relation to the zero divisor. The word independently refers to the mathematical construction: positivity of the Weil pairing, or a theorem already equivalent to it, cannot be an unnamed input.

Part I fixes the centered Mellin convention and examines geometric sufficient conditions [1]. Part II distinguishes finite prime formulas, spectral truncations, positive compressions, and ground-state approximations [2]. Here we examine what these routes have in common and what their comparison maps actually establish. The results in this paper are elementary structural reductions and counterexamples, rather than a construction of the required arithmetic geometry.

The slogan connecting purity, positivity, polarisation, and spectral realization is useful as a list of questions. It is not an existing chain of unconditional equivalences. Each term has several mathematical meanings. Purity can refer to weights of an arithmetic cohomology theory, the location of operator eigenvalues, or a predicate in a formal interface. Polarisation can mean a positive Hermitian form, a compatible complex structure on a fixed symplectic space, or a bilinear relation in a category. The maps between these meanings are part of the problem.

A proposition-level equivalence carries less information than a realization map. If two predicates are each equivalent to RH, composing the equivalences proves they are equivalent to each other. This supplies no vector space, no pairing, no transformation of test functions, and no estimate. A structural explanation must identify those objects and maps. It should say, for example, that an arithmetic cycle map pulls a negative intersection pairing back to the negative of the Weil pairing.

1.1 The fixed analytic target

We use the test algebra $\mathcal{T} = C_c^\infty(\mathbb{R}, \mathbb{C})$, centered transform F_f , involution f^* , and Hermitian form B_W defined in Part I. The convention is linear in the first argument. Thus the quadratic functional is $Q_W(f) = B_W(f, f)$. The underlying real vector space still contains every complex test function; realifying the scalar structure does not restrict the quantifier to real-valued tests.

For reference, the decisive distinction in the zero-side expression is

$$Q_W(f) = \sum_{\rho} F_f(\rho - \tfrac{1}{2}) \overline{F_f(\tfrac{1}{2} - \bar{\rho})}. \quad (1)$$

Zeros are counted with multiplicity. A summand becomes an absolute square when the associated zero lies on the critical line. Replacing the second factor by the conjugate of the first before establishing that fact would assume the desired conclusion. The smoothness and compact support of f give the decay needed for absolute convergence, with the explicit-formula convention of [3, Appendix B]. Convergence of this sum is consequently distinct from its nonnegativity.

The classical Weil criterion identifies nonnegativity of this exact functional on all of $\mathcal{T}$ with RH. It is a cited analytic theorem in this program, not a new theorem or an axiom introduced into the Lean library. The point of the formal work is to verify smaller transfers while exposing exactly where the arithmetic identification must enter.

1.2 Requirements for a structural explanation

A useful explanation should specify a source category, a realization functor, a positive or negative form in that category, and the comparison with the analytic target. It should also distinguish statements that are already known in that category from the construction still required for the arithmetic example. A Hodge index theorem on an available smooth projective surface is known mathematics. Producing an appropriate object associated to the arithmetic base, with the right correspondence algebra and archimedean terms, is a different task.

This separation prevents two common errors. One is to put the target identity into an opaque proposition named geometric realization and treat the name as a construction. The other is to demand so much structure that every later proof becomes formal, while leaving the existence

of the structure unexamined. A sufficient package is useful when its fields identify independent construction problems. It becomes circular when a field is the original positivity statement with renamed variables.

2 Positive realizations and relative minimality

The term minimal requires a comparison relation. There is no category-free smallest arithmetic-geometric object. We first fix a modest category in which minimality has an exact meaning. The result explains what a universal positivity object can provide and why that result alone does not settle the arithmetic question.

Throughout this section, V is a complex vector space and B is a Hermitian sesquilinear form, linear in its first argument. No positivity is initially assumed. A positive realization of (V, B) consists of a complex Hilbert space H and a linear map $j : V \rightarrow H$ such that

$$B(x, y) = \langle jx, jy \rangle_H \quad (x, y \in V). \quad (2)$$

The Hilbert inner product in this display is linear in the first argument. This convention differs from Lean's complex inner product; the formal spectral statements below explicitly use Lean's convention.

Definition 2.1. A morphism from a positive realization (H, j) to (K, k) is a linear isometric embedding $\Phi : H \rightarrow K$ satisfying $\Phi j = k$. A realization is generated if $j(V)$ is dense in H .

Composition is ordinary composition of linear maps. The commuting triangle is part of the morphism definition. It prevents an arbitrary isometry between unrelated Hilbert spaces from being counted as an arithmetic comparison.

Lemma 2.2. *If B is nonnegative, its diagonal null set equals its radical:*

$$\{x : B(x, x) = 0\} = \{x : B(x, y) = 0 \text{ for every } y\}.$$

This set is a linear subspace.

Proof. Suppose $B(x, x) = 0$ and choose y . Put $b = B(x, y)$ and $c = B(y, y) \geq 0$. Positivity of $B(x + ty, x + ty)$ for every complex t gives

$$0 \leq 2 \operatorname{Re}(\bar{t}b) + |t|^2 c.$$

If $b \neq 0$, choose $t = -\epsilon b$ with $\epsilon > 0$. The right side is $-2\epsilon|b|^2 + \epsilon^2|b|^2 c$, which is negative for sufficiently small ϵ . Hence $b = 0$. The converse follows by setting $y = x$. The radical is a subspace by sesquilinearity. $\square$

The same argument works over the reals using real scalars. The Lean quotient construction uses the real bilinear version. Nonnegativity is essential: for an indefinite form an isotropic vector need not be orthogonal to the entire space, and its diagonal null set need not be a subspace.

Theorem 2.3. *A positive realization of (V, B) exists if and only if B is nonnegative. When it exists, there is an initial positive realization (H_B, j_B) . It is generated. Every generated positive realization is uniquely isometrically isomorphic to it, with the isomorphism respecting the maps from V .*

Proof. A realization makes $B(x, x)$ a squared norm and therefore nonnegative. Conversely suppose B is nonnegative and let $N = \text{rad } B$. The formula

$$\langle [x], [y] \rangle = B(x, y)$$

defines a positive definite inner product on V/N by Lemma 2.2. Complete this inner product space to H_B , and let j_B be the quotient map followed by the completion embedding.

Given any positive realization (K, k) , a vector x belongs to N precisely when $\|kx\|^2 = B(x, x) = 0$. Thus $[x] \mapsto kx$ is well-defined and isometric on V/N . An isometry into a complete space extends uniquely to the completion. Its extension $\Phi : H_B \rightarrow K$ satisfies $\Phi j_B = k$ and remains an isometry by continuity of the norm. Uniqueness follows from density of $j_B(V)$. If $k(V)$ is dense, the image of Φ is both dense and closed, so Φ is onto. This proves the final assertion. $\square$

This is the familiar quotient-and-completion mechanism underlying positive kernel realizations [4]. It is included with proof because its exact logical role matters here. The construction begins after nonnegativity has been established. Applying it to $(\mathcal{T}, B_W)$ gives a universal Hilbert realization only under Weil positivity. Its universal property does not supply positivity for free.

Corollary 2.4. *For the exact Weil pairing, existence of the positive realization in Definition 2.1 is equivalent to Weil positivity, and therefore to RH by the classical criterion.*

This corollary is a circularity test, not a proposed route to prove RH. A candidate construction of H_B must be inspected to determine where it established nonnegativity without invoking this corollary backwards. Calling H_B a cohomology space or a universal arithmetic Hilbert space does not change its prerequisites.

2.1 Weaker certificates and their limits

If one asks only for a function $r : V \rightarrow \mathbb{R}$ with $B(x, x) = r(x)^2$, existence is again equivalent to nonnegativity: take the nonnegative square root pointwise. This certificate forgets addition, scalar multiplication, and the off-diagonal form. It is logically sufficient for the sign but structurally poor. The Lean library proves this equivalence explicitly so that an arbitrary norm-square representation is not mistaken for a new arithmetic construction.

Linearity of a realization restores the off-diagonal information through polarization. For a Hermitian form, the diagonal determines the form because the four values at $x + y$, $x - y$, $x + iy$, and $x - iy$ recover the real and imaginary cross terms. Consequently a linear map preserving every diagonal value preserves the entire form. This is a genuine additional property of the map, even though existence of a linear Hilbert realization remains equivalent to positivity.

A comparison order can be defined by forgetting structure. A geometric certificate may forget its cycle algebra, degrees, and pullback operations to become a positive realization. The existence of that forgetful map does not imply that every positive realization comes from geometry. Minimality in the category of Hilbert realizations therefore says nothing about initiality in a category of arithmetic spaces.

3 The geometric comparison as a pullback

The primitive intersection package of Part I supplies a direct route to B_W . Its data include a correspondence space C , a Hermitian intersection form I , degree maps whose common kernel is

a primitive subspace P , and a linear map $\Gamma : \mathcal{T} \rightarrow P$. The two substantive assertions are the geometric inequality on P and the arithmetic comparison

$$B_W(f, g) = -I(\Gamma f, \Gamma g). \quad (3)$$

The sign in this formula is part of the data. An intersection form with the wrong sign does not produce Weil positivity.

Proposition 3.1. *Suppose $I(v, v) \leq 0$ for every $v \in P$, the map Γ is defined on all of $\mathcal{T}$, and equation 3 holds for every pair of tests. Then $Q_W(f) \geq 0$ for every $f \in \mathcal{T}$.*

Proof. For each f , the vector Γf belongs to P . The right side of equation 3, evaluated at (f, f) , is the negative of a nonpositive real number. $\square$

Neither continuity nor injectivity of Γ is needed for this sign transfer. Those hypotheses become relevant for extending the map, comparing radicals, or proving strict positivity. This distinction matters when describing minimal sufficient data. Requiring a complete Hilbert realization at this stage adds obligations that the elementary implication does not use.

The geometric content lies upstream of this proposition. The ordinary scheme $\text{Spec } \mathbb{Z}$ has dimension one. The surface argument for a curve over a finite field uses a square and an intersection theory of correspondences; the same construction has not been supplied merely by writing a square of the arithmetic base. A proposed arithmetic replacement must explain its base, the infinity contribution, the degree directions, the involution on correspondences, and why its trace is the specified completed-zeta distribution. Part I examines available geometric programs and keeps these construction questions separate from the Hodge transfer.

3.1 Logical strength of the image-restricted inequality

Suppose equation 3 has been proved, but the only proposed negativity theorem is

$$I(\Gamma f, \Gamma f) \leq 0 \quad (f \in \mathcal{T}).$$

This statement is exactly Weil positivity after substitution. It can still be a useful reformulation, because the images Γf may admit geometric tools. It cannot be called an already established geometric input without a proof using those tools.

A negativity theorem for an independently constructed primitive space is stronger data. Its usefulness comes from the independent construction and its proof, not simply from the fact that the theorem quantifies over more vectors. One could always define P using the Weil form and restate the same problem there. The source of the objects must therefore be recorded alongside the theorem signatures.

3.2 Composition and descent

Suppose another space P' maps to P by a linear map u preserving intersection forms, and $\Gamma = u\Gamma'$ for a map $\Gamma' : \mathcal{T} \rightarrow P'$. Then equation 3 pulls back along u . This is the basic compositional property APP needs: exact comparisons compose, and positivity pulls back along them.

If a pairing descends through a quotient, vectors killed by the quotient must pair to zero with every vector. For a nonnegative symmetric form the null-radical lemma provides precisely that condition. Thus the quotient construction can be formalized independently once positivity is given. It does not create a sign from an indefinite form. In a categorical setting, descent of the underlying

object and descent of a positive form are different obligations; the compatibility equations must include the form and any operators that act on it.

For genuinely local geometric data, a descent assertion would have to specify the covering family, overlap maps, cocycle equations, and the topology used for completion. None of these is forced by the words condensed or cohomological. A general sheaf or condensed category may provide a useful home for analytic objects without proving the arithmetic comparison or the sign of a pairing in that category.

4 A spectral bridge with an explicit domain

There is a short direct implication from a positive inner product and an adjoint weight identity to the real parts of eigenvalues. Its simplicity makes it a useful place to expose the otherwise hidden arithmetic assumptions.

Let H be a complex inner product space, with the inner product conjugate-linear in the first variable in this section. This convention matches `mathlib` and the Lean spectral statements below. Let $D \subset H$ be a complex linear subspace and $\Theta : D \rightarrow H$ a linear operator. No boundedness or self-adjoint extension is assumed. Fix a real weight w and suppose

$$\langle \Theta x, x \rangle + \langle x, \Theta x \rangle = w \langle x, x \rangle \quad (x \in D). \quad (4)$$

A full adjoint relation for pairs (x, y) is stronger than this diagonal identity. The weaker assertion suffices for the following theorem.

Theorem 4.1. *If $x \in D$ is nonzero and $\Theta x = zx$, then $\operatorname{Re} z = w/2$.*

Proof. Substitution in equation 4 gives

$$(\bar{z} + z) \|x\|^2 = w \|x\|^2.$$

The squared norm is a strictly positive real number. Taking real parts and dividing gives $2 \operatorname{Re} z = w$. $\square$

The theorem is implemented in `APP/Spectral.lean`, with D a submodule and Θ a linear map out of that submodule. Completeness is absent because the proof does not use it. The operator domain, eigenvector equation, and nonzero vector hypothesis all appear explicitly.

Definition 4.2. For an indexed family $z_i \in \mathbb{C}$, an eigen-realization in (D, Θ) means that for every i there exists a nonzero $x_i \in D$ with $\Theta x_i = z_i x_i$.

Combining an eigen-realization with Theorem 4.1 forces every z_i onto the line $\operatorname{Re} z = w/2$. For $w = 1$, that line is the critical line. For an arithmetic application, however, the indexed family must be the actual nontrivial zeros, with the required multiplicities if a trace or determinant identity is claimed.

A self-adjoint operator with a chosen real spectrum always satisfies an appropriate reality theorem. That observation does not identify its spectrum with the zeros of zeta. Counting asymptotics, low-lying numerical agreement, or the availability of a trace formula with omitted terms are insufficient replacements for the realization hypothesis. They provide evidence or partial comparisons, not the nonzero eigenvectors required by Definition 4.2.

4.1 From purity to an invariant metric

A converse needs more care. In finite dimensions, a diagonalizable complex operator whose eigenvalues all have modulus one admits a positive invariant Hermitian form: choose an eigenbasis and make it orthonormal. The same conclusion fails without diagonalizability. It also does not extend verbatim to infinite-dimensional representations, where the choice of a complete norm and boundedness of the action matter.

Example 4.3. On $\mathbb{R}^2$, let $S(x, y) = (x + y, y)$. Every eigenvalue of the complexified operator is 1, but there is no positive definite bilinear form invariant under S .

Proof. The matrix of S is upper triangular with diagonal entries 1. Let $e_1 = (1, 0)$ and $e_2 = (0, 1)$. If a bilinear form G were invariant, then

$$G(e_1, e_2) = G(Se_1, Se_2) = G(e_1, e_1 + e_2) = G(e_1, e_1) + G(e_1, e_2).$$

Thus $G(e_1, e_1) = 0$, contradicting positive definiteness. $\square$

This example refutes a proposed bridge from bare spectral purity to invariant polarisation. It does not refute stronger cohomological purity packages that include semisimplicity or a polarization theorem. Rather, it identifies which additional structure such a bridge must use. The Lean library verifies the shear obstruction in ordinary real linear algebra.

5 Polarisation from a positive Weil realization

Once a positive Hilbert realization has been constructed, it has a canonical real symplectic structure. This is a useful direct bridge, provided its direction is stated correctly. Continue with the first-variable-linear convention on a complex Hilbert space H . On its underlying real space set

$$\omega(x, y) = -\operatorname{Im}\langle x, y \rangle, \quad Jx = ix, \quad g(x, y) = \operatorname{Re}\langle x, y \rangle. \quad (5)$$

Then ω is real bilinear and skew, $J^2 = -1$, and

$$\omega(Jx, Jy) = \omega(x, y), \quad \omega(x, Jy) = g(x, y), \quad g(x, x) = \|x\|^2.$$

Thus J is a positive compatible complex structure for this particular ω . It commutes with every complex-linear unitary action on H .

The sign in equation 5 follows from the inner product convention. Indeed, $\langle x, iy \rangle = -i\langle x, y \rangle$, so its imaginary part is $-\operatorname{Re}\langle x, y \rangle$. This calculation prevents a sign convention from being confused with a positivity theorem.

Applied to the universal realization of the Weil form, this construction gives a positive polarisation after Weil positivity. It does not establish a polarisation of an independently fixed wedge object. To make that claim, one must provide a real linear comparison with that object, prove that it identifies the skew forms, and show that the relevant actions are intertwined.

Proposition 5.1. *Let (E, ω_E) be a real vector space with a skew form. Suppose $U : E \rightarrow H$ is an injective real linear map, its image is stable under J , and*

$$\omega_E(x, y) = \omega(Ux, Uy).$$

There is a unique real linear operator J_E with $UJ_E = JU$. It satisfies $J_E^2 = -1$, preserves ω_E , and obeys

$$\omega_E(x, J_E x) = \|Ux\|^2 > 0 \quad \text{for } x \neq 0.$$

If an action on E is intertwined with a complex-linear action on H , then J_E commutes with it.

Proof. Image stability defines $J_E = U^{-1}JU$, where the inverse is taken on $U(E)$. Linearity and uniqueness follow from injectivity. Applying U to J_E^2x gives $-Ux$, so $J_E^2x = -x$. Preservation and positivity follow from the displayed comparison and the identities on H . For an action element R and its complex-linear counterpart $\tilde{R}$, the identities $UR = \tilde{R}U$ and $\tilde{R}J = J\tilde{R}$ imply $URJ_E = UJ_ER$. Cancel U . $\square$

This algebraic transport does not assert that J_E is continuous in a pre-existing norm on E . It also does not assert that $g_E(x, y) = \omega_E(x, J_E y)$ defines a complete norm. If U is an isomorphism of the intended topological spaces, the corresponding continuity statements follow. If it is merely a dense injection, completeness requires additional work.

5.1 The radical and the intended carrier

The universal realization factors through the radical of the Weil form. A comparison with a wedge space must therefore explain whether that wedge space represents the full test space, its radical quotient, a completion, or another representation entirely. These choices are mathematically different even when their names suggest the same object.

A zero vector space has a vector, namely zero. It is not an empty type. Thus an assertion that a linear defect vanishes is ordinarily expressed by the space being the zero space, or by its nonzero-vector subtype being empty. Replacing this by emptiness of the full vector carrier changes the statement. This distinction already arises before any analytic positivity argument and is relevant to the context-repository audit in Part I.

An exact comparison with a positive target also identifies the source radical with the kernel of the realization map. Injectivity is therefore a substantive hypothesis when the desired conclusion is strict positivity on the source. A source vector killed by the map cannot acquire a strictly positive norm by pulling back the target metric. Quotienting it out may be correct, but the quotient has to match the intended arithmetic object.

6 Limits of certificates on one test space

Finite constructions can support a limiting argument only after they have been compared on common inputs. A family of positive spaces H_n does not by itself define a limit of forms on $\mathcal{T}$. Maps $d_n : \mathcal{T} \rightarrow H_n$ are needed, or compatible maps from increasing subspaces that eventually contain each test.

Let q_n be nonnegative quadratic functionals on H_n . Set

$$q_n^{\mathcal{T}}(f) = q_n(d_n f).$$

The elementary limit mechanism used in Part II is that, for every fixed f , convergence $q_n^{\mathcal{T}}(f) \rightarrow Q_W(f)$ implies nonnegativity of $Q_W(f)$. No operator norm convergence is needed for this conclusion.

Proposition 6.1. *Suppose that for each $f \in \mathcal{T}$ there are real numbers $\epsilon_n(f) \rightarrow 0$ and, eventually in n , positive-model values $q_n(d_n f) \geq 0$ satisfying*

$$Q_W(f) \geq q_n(d_n f) - \epsilon_n(f).$$

Then $Q_W(f) \geq 0$ for every f .

Proof. Fix f . The inequality gives $Q_W(f) \geq -\epsilon_n(f)$ eventually. The right side tends to zero. If $Q_W(f) < 0$, the inequality fails for all sufficiently large n , a contradiction. $\square$

This one-sided defect estimate is weaker than convergence of the full form difference in absolute value. It identifies a smaller sufficient analytic task when a positive finite model is independently available. In particular, one need not prove convergence of the vectors $d_n f$ in an unspecified ambient Hilbert space to apply this proposition.

There is also a simpler formulation that no longer requires a model: $Q_W(f) \geq -\epsilon_n(f)$ with $\epsilon_n(f) \rightarrow 0$. That formulation is equivalent to positivity, by taking every error to be zero in the converse direction. The positive finite model matters because it supplies an independently analyzable source for the lower bound. Merely writing an error sequence does not establish it.

6.1 The quantifier order

For the sign transfer, the needed order is

$$\forall f \in \mathcal{T} \quad \forall \epsilon > 0 \quad \exists N \quad \forall n \geq N : \quad Q_W(f) \geq -\epsilon.$$

A common N for every unit vector is unnecessary. Conversely, a calculation for finitely many chosen tests supplies only those finitely many inequalities. Increasing the number of tests in numerical experiments does not replace the all-test quantifier unless a density theorem and error estimate are proved.

One may work on a dense common core when the target quadratic form is continuous in the chosen topology. Here the topology must be named. Smooth compact tests have a natural test-function topology; an unbounded quadratic form on a Hilbert space is generally not continuous in the ambient Hilbert norm. A statement about a form core uses the form norm, not merely the norm in the surrounding L^2 space.

Finite-dimensional positivity is also basis-dependent in its interpretation. A positive semidefinite matrix representing the exact form on a finite span proves positivity on that span. A sampled kernel matrix may instead represent a quadrature approximation. In that case its positivity concerns the discretized form, and a quadrature error estimate is another comparison obligation. These are different applications of the same pullback principle.

6.2 Subtraction and normalization

A positive operator has a nonnegative ordinary trace whenever that trace is defined. A renormalized expression is often a difference of divergent or cutoff-dependent positive quantities. The sign of the difference does not follow from the signs of its summands.

Example 6.2. For $x \in \mathbb{R}$, let $p_n(x) = nx^2$ and $c_n(x) = (n+1)x^2$. Both are nonnegative, but $p_n(x) - c_n(x) = -x^2$ for every n .

This example applies to the logical form of a subtraction argument, not to a particular trace formula. It shows which inference needs justification. In the arithmetic explicit formula, the prime-power contribution and the archimedean subtraction have fixed signs and normalizations. A proposed positive compression must be compared with that entire expression, including all correction terms.

A related issue arises from shifting a lower-bounded self-adjoint operator, as in the localized constructions discussed by Suzuki [5]. If A has lower bound m , then $A - cI$ is positive for $c < m$. This gives a useful Hilbert norm, but it says nothing about positivity of A when c is negative. The scalar operator $A = -I$ becomes positive after subtracting $-2I$. Returning to the unshifted arithmetic form is a separate step.

7 The ground-state route and its smaller residual

The Connes finite-approximation program has a more specific structure than an arbitrary positive-form limit. It relates a localized Weil form to a ground-state eigenfunction, associates real zeros to an appropriate Fourier transform, and compares an explicit prolate-based proxy with the Riemann Ξ -function. The primary sources distinguish the proxy limit from the missing comparison with the actual ground state [6, 7]. The precise real-zero theorem also requires the stated distributional operator and trigonometric core to satisfy its essential self-adjointness hypotheses, as well as a simple isolated minimum and an even eigenfunction [8, Theorem 6.1].

Let a be a logarithmic support parameter, η_a an actual ground-state eigenfunction, and k_a a specified proxy, both written in additive coordinates. The source normalization of the proxy is chosen so that its transform tends to a fixed nonzero multiple of Ξ . Absorb that fixed constant into the notation. A sufficient remaining comparison is the existence of nonzero scalars c_a for which

$$\sup_{z \in K} |c_a \widehat{\eta_a}(z) - \widehat{k_a}(z)| \longrightarrow 0 \quad (6)$$

for every compact K in the relevant open strip. The constants c_a must be fixed consistently in the limit, rather than chosen separately at each evaluation point.

7.1 An explicit sufficient norm estimate

For a closed substrip $|\operatorname{Im} z| \leq b$, a stronger but concrete estimate is

$$\int_{\mathbb{R}} |c_a \eta_a(t) - k_a(t)| e^{b|t|} dt \longrightarrow 0. \quad (7)$$

Indeed, $|e^{izt}| \leq e^{b|t|}$ on that substrip. Taking absolute values inside the Fourier integral bounds the left side of equation 6 by the integral in equation 7. This estimate supplies a direct map from an eigenfunction comparison to a holomorphic transform comparison.

If the difference is supported in $[-a, a]$, Cauchy-Schwarz gives the cruder sufficient bound

$$\int_{-a}^a |c_a \eta_a - k_a| e^{b|t|} dt \leq (2a)^{1/2} e^{ba} \|c_a \eta_a - k_a\|_{L^2}.$$

Therefore unweighted L^2 convergence with no rate is insufficient for this particular argument when $a \rightarrow \infty$. One can seek a sharper weighted estimate, a different transform bound, or a compactness argument with a normal-family hypothesis. What one cannot do is omit the growing weight while claiming uniform convergence on complex compact sets.

7.2 Rayleigh bounds and ground-state selection

Suppose an approximate vector has a small value of the Weil form. The ground-state Rayleigh quotient is no larger than that value, by the variational principle. This gives an upper bound for the lowest eigenvalue, whereas positivity requires a lower bound. A small trial value thus has the wrong inequality direction for proving positivity by itself.

Similarly, a small residual in an eigenvalue equation does not determine which eigenspace contains the approximate vector. Even after the intended eigenvalue has been correctly selected, a shrinking spectral gap can prevent residual convergence from controlling the vector.

Example 7.1. Let A_n on $\mathbb{R}^2$ have matrix $\operatorname{diag}(0, (n+1)^{-1})$, and let $v = e_2$. The ground eigenspace is $\mathbb{R}e_1$, while

$$\|A_n v\| = (n+1)^{-1} \longrightarrow 0, \quad \operatorname{dist}(v, \mathbb{R}e_1) = 1.$$

For a self-adjoint operator with a simple ground eigenvalue λ_0 separated from the rest of its spectrum by $\delta > 0$, the spectral theorem instead gives

$$\|(1 - P_0)v\| \leq \delta^{-1}\|(A - \lambda_0)v\|.$$

The inequality follows by integrating $(t - \lambda_0)^2 \geq \delta^2$ on the complement of the ground spectral projection. In a varying family, the useful condition is control of the residual divided by the gap, together with phase and normalization. Neither a vanishing residual alone nor a plausible trial eigenvalue supplies that condition.

7.3 The holomorphic zero-transfer obligation

Assume the transforms $c_a \hat{\eta}_a$ are holomorphic on a connected open strip and have no zeros off the real axis. Suppose they converge uniformly on compact subsets to Ξ , which is not identically zero. Hurwitz's theorem applied on each connected off-real component makes the limit either nowhere zero or identically zero there. The second possibility would make Ξ vanish on an open subset of the connected strip, and hence on the whole strip by the identity theorem, contrary to the nonzero-limit hypothesis. Thus Ξ has no zeros off the real axis in the strip. This is the precise complex-analytic transfer needed for the real-zero route.

The theorem requires holomorphic convergence, not agreement of sampled real values. Multiplication by a nonzero scalar does not change zeros, whereas a variable normalization must be holomorphic and nowhere zero on the domain. A meromorphic ratio requires restriction to domains avoiding its poles, and cannot be inserted into this argument as if it were an entire limit. These conditions become material when comparing alternative spectral determinant formulations.

The smaller residual for this route has two separate parts: the hypotheses ensuring the real-zero property for the actual ground-state transform throughout an exhausting family of intervals, and the normalized comparison in equation 6. General Hurwitz theory is already available. The missing theorem is the arithmetic and analytic identification of the family to which it is applied.

8 A network of bridges with typed obligations

A useful diagram for APP is a network of sufficient maps, not a cycle of unexplained equivalences. It begins with several candidate source structures and records what additional data reach the common analytic pairing.

Source data	Required comparison	Consequence
Primitive geometric pairing	Signed pullback along Γ	Weil nonnegativity
Positive Hilbert realization	Exact pairing on all tests	Weil nonnegativity
Weight identity on a domain	Nonzero eigenvectors for actual zeros	Critical-line locations
Positive finite models	Common tests and vanishing lower defect	Weil nonnegativity
Ground-state transforms	Real-zero hypotheses and normalized holomorphic limit	Real zeros of Ξ
Positive Weil quotient	Skew-form and action identification	Polarisation on the identified object

The reverse direction in a row is generally unavailable. The universal Hilbert realization gives an equivalence only after the target form is fixed. The geometric source in the first row carries additional arithmetic meaning. The last row concerns an identified target object; it is not a construction of every wedge object appearing in the context repositories.

A diagram becomes commutative through equalities of maps or pairings. For example, if two geometric realizations Γ_1 and Γ_2 are related by an isometry u with $u\Gamma_1 = \Gamma_2$, their pulled-back forms agree. If two pipelines separately conclude RH, that conclusion alone does not supply u . The distinction is the mathematical reason to formalize the bridge maps themselves.

8.1 A disciplined meaning of the common obstruction

The common obstruction is not a newly proved universal object. It is the failure, so far, to construct a positive arithmetic realization together with an exact comparison to the Weil pairing. Each route makes a different part of that problem concrete. Geometry asks for the object, its primitive sign theorem, and its trace comparison. Finite positive models ask for a lower-defect bound on all tests. The ground-state route asks for selection and transform convergence. The spectral route asks for an actual zero realization in a positive space with the correct weight identity.

These are not automatically equivalent existence problems in their stated categories. Some packages may be strictly stronger than Weil positivity. Others, such as an existential Hilbert factorization of the exact form, are equivalent by construction. The comparison order and the source of each assumption must be included before making a minimality claim.

The user of a formal library should be able to locate the arithmetic obligation without unfolding an opaque proposition. In APP, the comparison fields are equalities or quantified eigenvector equations. No target theorem is declared as an axiom. The mathematical library verifies that these fields suffice, and the research notes record which concrete arithmetic instances remain absent.

9 Formal coverage and unproved arithmetic instances

The Lean development separates the universal algebra from the analytic target. `APP/Structures.lean` defines nonnegativity, exact arithmetic comparisons, norm-square certificates, and primitive intersection comparisons. `APP/Bridges.lean` handles radicals, quotient forms, and failures of overly broad converse claims. `APP/Approximation.lean` proves sign transfer under explicitly stated limiting hypotheses. `APP/Spectral.lean` implements the domain-level weight argument. `APP/Residual.lean` isolates an eventual one-sided comparison defect and proves that unrestricted existential model certificates remain equivalent to target positivity.

The exact classical Weil distribution, its explicit formula, and the equivalence with RH are not yet formalized in this library. They are defined and sourced in the mathematical papers. The Lean transfer results are parameterized by an actual real-valued function or bilinear form, so the absent arithmetic instantiation remains visible. This is a useful formal boundary: Lean verifies the implication, but the library does not claim to have supplied its number-theoretic arguments.

9.1 Declarations versus theorem instances

A theorem parameterized by Q and a comparison map proves a general implication for every valid instance. An arithmetic result requires a term of the relevant structure with $Q = Q_W$. The existence of the structure declaration is not evidence that such a term has been constructed. The distinction applies equally to a cohomology interface, a positive-kernel interface, and a convergence certificate.

The elementary nonzero examples in the library serve a narrower purpose. They establish that the interfaces have ordinary mathematical instances and that the transfer lemmas are not formulated with contradictory hypotheses. They do not test the arithmetic instance. A two-dimensional primitive form is an example of the mechanism, not a surrogate arithmetic surface.

The axiom audit records kernel dependencies of exported results. Standard logical axioms used by Lean and mathlib are distinguished from mathematical assumptions carried as explicit theorem parameters. A clean axiom audit says that the Lean proofs contain no newly postulated arithmetic theorem. It cannot establish that the user-facing identification of an abstract Q with the classical Weil form is correct; that identification must be implemented and reviewed separately.

9.2 The next independently testable questions

For a geometric proposal, the first concrete questions concern the correspondence algebra and the exact arithmetic pairing. An independently proved negative intersection theorem is useful only after its primitive subspace and test-function map have been specified. At infinity, the chosen regularization must reproduce the completed-zeta term, including the subtraction at the identity.

For finite models, the first questions concern the finite form actually computed, the map from smooth tests to the model, and a one-sided comparison bound. It is often better to estimate a single scalar defect on a fixed test than to seek operator norm convergence on an entire ambient unit ball. This choice reduces the analytic burden without changing the all-test quantifier required at the end.

For the ground-state route, an informative calculation must distinguish the proxy, the exact ground state, and the gap above it. A rigorous bound on their normalized difference would connect two constructions that are presently related by evidence and conjecture. A small value of the quadratic form is not that bound. Nor does exact truncation of the prime sum identify the extremizing vector.

For the cohomological route, the first questions concern a nontrivial realization and a genuine weight identity. A predicate declaring a purity bound without referring to an action or eigenvalues does not answer them. The context audit shows why replacing such predicates by maps and identities is a substantive reformulation even when it does not yet prove the arithmetic theorem.

10 Finite Gram data and compatible limits

Finite positivity becomes structurally meaningful when the finite objects are related by maps preserving their forms. This is a different condition from merely producing a positive matrix at each cutoff. The distinction can be expressed without any reference to zeta and then applied to a specified arithmetic dictionary.

Let S be a set of labels and $K : S \times S \rightarrow \mathbb{C}$ a Hermitian kernel, linear coefficients being placed in the first slot of its quadratic expression. For a finite list $s_1, \dots, s_m$ and coefficients $c_1, \dots, c_m$, define

$$q_K(c; s) = \sum_{i,j=1}^m c_i \overline{c_j} K(s_i, s_j).$$

The kernel is of positive type if this real number is nonnegative for every list and every coefficient vector. Repeated labels are allowed. This definition records all finite linear combinations; checking only the diagonal $K(s, s)$ is not enough.

Example 10.1. The Hermitian matrix

$$\begin{pmatrix} 1 & 2 \\ 2 & 1 \end{pmatrix}$$

has positive diagonal entries. At the coefficient vector $(1, -1)$ its quadratic value is -2 . Consequently positive self-pairings of dictionary entries do not imply positivity on their span.

If $K(s, t) = B(v_s, v_t)$ for vectors in a test space, the Gram expression equals $B(\sum c_i v_{s_i}, \sum c_i v_{s_i})$ by sesquilinearity. This exact identity is the finite-dimensional comparison map. If the vectors span the entire test space, positivity of every such Gram matrix is equivalent to positivity of the form. If they only span a dense subspace, extension requires continuity in the specified topology.

Proposition 10.2. *Suppose K_n are positive-type kernels on one label set S and $K_n(s, t) \rightarrow K(s, t)$ for every pair. Then K is of positive type.*

Proof. Fix a finite list and a coefficient vector. Its Gram expression is a finite linear combination of entries of K_n , so it converges to the corresponding expression for K . Each expression for K_n is nonnegative. The closedness of the nonnegative real half-line gives the result. $\square$

Only finitely many limits are involved after the list is fixed. No uniformity in S is necessary for this conclusion. The Lean library includes the real finite-matrix version. An application to a complex Hermitian kernel may use its realification on the full complex coefficient space, provided that realification is made explicit.

10.1 Consistent finite kernels

Suppose the label set is an increasing union $S = \bigcup_n S_n$ of finite sets, and each K_n is a kernel on S_n . If

$$K_{n+1}|_{S_n \times S_n} = K_n,$$

there is a well-defined kernel K on S : choose an index containing both labels and use that value. The equality above makes the value independent of the choice. If each finite kernel is positive, every finite list lies in some S_n , and K is positive.

This construction is exact descent of finite Gram data. It supplies the inner-product quotient and its completion by Theorem 2.3, applied to the free vector space on S . The construction is independent of RH as a general theorem. For an arithmetic application, however, one must prove that the resulting kernel is the kernel obtained from B_W on the intended dictionary. Without that comparison, one has constructed a positive Hilbert space associated to other data.

A merely approximate consistency condition is more delicate. For each fixed pair of labels, a summable bound on successive entry differences makes the entries Cauchy. Positivity then passes to the resulting kernel by Proposition 10.2. This is an alternative intermediate structure: compatible finite Gram kernels can construct the positive target space rather than requiring a pre-existing common Hilbert space. The independent arithmetic comparison remains essential.

10.2 Direct limits of inner product spaces

The same idea can be expressed with finite-dimensional inner product spaces H_n and isometric embeddings $u_n : H_n \rightarrow H_{n+1}$. A vector in the algebraic direct limit is represented by a pair (n, x) ; two representatives agree if their images agree at a later stage. The inner product of two classes is computed after mapping both representatives into a common later space. Isometry makes the result independent of the stage.

Proposition 10.3. *A sequential system of inner product spaces and isometric linear embeddings has an algebraic direct limit with a positive definite inner product. Its Hilbert completion has the expected isometric comparison maps. Every compatible family of isometries into a fixed Hilbert space factors uniquely through this completion.*

Proof. The inner product is well-defined by preservation at later stages. If a representative has zero squared norm in the limit, it has zero squared norm at its original stage and hence is zero there. This proves definiteness. A compatible family of isometries assigns the same target vector to equivalent representatives, so it defines an isometry on the algebraic limit. Complete that isometry as in Theorem 2.3. Density gives uniqueness. $\square$

This proposition describes a possible stronger intermediate structure for finite approximation programs. It includes actual transition maps and exact preservation of the pairing. A collection of finite self-adjoint matrices with no maps between their underlying spaces does not provide this system. Nor does convergence of a few eigenvalues specify transition maps preserving all cross terms.

10.3 Operators on the compatible limit

Suppose operators J_n on H_n commute with the transition maps, square to -1 , and preserve the inner products. Then they define a linear isometry J on the algebraic direct limit. In particular $\|Jx - Jy\| = \|x - y\|$, so J is uniformly continuous and extends uniquely to the completion. The extension is linear and isometric, and satisfies $J^2 = -1$ there by continuity. An action compatible at every stage and commuting with every J_n also commutes with the limit operator, provided the action extends continuously.

This construction is a genuine descent mechanism for a polarisation already present at every stage. Its hypotheses are stronger than positivity of a scalar limit, because they prescribe compatible vectors and operators. That extra structure can be valuable: it controls the realized object, not just the sign of an arithmetic functional. It must nevertheless be supplied by the concrete approximation family. An isometry between successive abstract Hilbert spaces may exist without intertwining the required arithmetic maps or operators.

The direction of investigation depends on the available data. If exact isometric transitions are natural, the direct-limit construction is appropriate. If only scalar form comparisons are accessible, the one-sided residual of Proposition 6.1 asks for less. APP does not identify these routes as equivalent; it records a forgetful passage from a compatible positive system to positive scalar comparisons whenever the arithmetic realization maps commute.

11 Arithmetic operator comparisons

A concrete comparison is more informative than an abstract existence criterion because its image can be analyzed. The localized screw-function formulation discussed in Part II provides such a case. On $(-a, a)$, differentiation sends a Dirichlet test function to a mean-zero function. This codomain restriction is a boundary condition, not an arbitrary deletion of troublesome vectors.

In the notation of Part II, $D = i d/dt$ maps $H_0^1(-a, a)$ onto $L_0^2(-a, a)$, the mean-zero subspace. For $u \in L_0^2(-a, a)$, the inverse is

$$(Ku)(t) = -i \int_{-a}^t u(s) ds.$$

Its values vanish at both endpoints because u has integral zero. Cauchy-Schwarz gives $\|Ku\|_2 \leq 2a\|u\|_2$, and its weak derivative is $-iu$. Thus $DKu = u$. The precise domains make the map surjective, which is what permits reflection of a positivity statement in addition to pullback.

Suppose the exact localized comparison reads

$$Q_W^a(v) = \langle G_a Dv, Dv \rangle \quad (v \in H_0^1(-a, a)).$$

Then positivity of G_a on its mean-zero space implies positivity of the localized Weil form on the stated domain. Conversely, if the form is nonnegative there, setting $v = Ku$ shows $\langle G_a u, u \rangle \geq 0$ for every mean-zero u . These directions use different facts: pullback uses the comparison equality, while reflection also uses surjectivity of D .

This example answers part of the request for maps explaining equivalences. The equivalence is obtained by an explicit differential map with an explicit inverse on a specified subspace. It does not pass through RH. It also shows where the unresolved sign remains: proving positivity of the concrete G_a for every interval is not an independently known input. It is a concrete reformulation with access to an integral kernel.

The source construction of a self-adjoint localized Weil operator is not the same as positivity of that operator. Lower boundedness permits a Friedrichs extension of the specified symmetric operator, and compactness properties can supply discrete spectrum. The lower spectral bound may still be negative. Likewise, the positivity of a shifted operator supplies a Hilbert norm at that shift without identifying the unshifted arithmetic form with a squared norm.

These distinctions suggest a productive next step: retain the exact differential comparison and study the sign through the kernel, or derive lower bounds that compare a finite positive kernel with G_a . The all-interval quantifier must remain visible. The finite nature of the prime sum at each interval does not turn that infinite family of inequalities into a single finite calculation.

12 Research outcome

The minimal sign-transfer mechanism is a pullback of a nonnegative form, or equivalently the negative of a primitive nonpositive intersection form. A Hilbert realization has a universal quotient-and-completion description once positivity is known. These facts organize the program but do not determine a smallest independently constructed arithmetic-geometric object over $\text{Spec } \mathbb{Z}$.

Finite approximations provide several genuine pieces: exact arithmetic evaluations at fixed support, convergent smooth zero sums, localized operator constructions, and particular proxy limits. They have not been shown here to converge to the full positive arithmetic realization. The remaining work is recorded as explicit comparison, selection, and sign-control obligations. Some are known general theorems awaiting an arithmetic instance; others are the unresolved arithmetic content itself.

This outcome supports a structural research program with falsifiable intermediate tasks. It does not support an RH proof claim. The value of the network is that every proposed arrow now has a domain, a codomain, and a property to preserve. An attempted construction can fail at one of those fields, where a new estimate or geometric theorem can then be sought.

A Small countermodels

The countermodels below are elementary enough to verify directly. Their purpose is to locate invalid implications, not to suggest that the arithmetic examples must exhibit the same pathology.

A.1 Dropping the comparison

Take $V = \mathbb{R}$ with target $Q(x) = -x^2$ and a positive model $H = \mathbb{R}$ with $q(y) = y^2$. The model exists and is complete. No map d can satisfy $Q(x) = q(dx)$ for every x , because the two sides have opposite signs at $x = 1$. Thus existence of a positive model says nothing about the sign of an unrelated arithmetic expression.

A.2 Dropping injectivity

Take a nonzero vector space V and the zero map into any positive inner product space. Its pulled-back quadratic form is identically zero and is nonnegative. It is not strictly positive on nonzero vectors. This is why a positivity certificate suffices for the Weil criterion but a positive definite metric on a specified source requires kernel information.

A.3 Dropping nonnegativity before quotienting

On $\mathbb{R}^2$ let $B((x_1, x_2), (y_1, y_2)) = x_1 y_1 - x_2 y_2$. The vector $(1, 1)$ has zero diagonal value but pairs to 1 with $(1, 0)$. Its class cannot be killed while retaining this pairing. The null cone consists of two lines and is not a subspace. Quotienting the null cone is therefore not a general construction for indefinite forms.

A.4 Dropping semisimplicity

The shear of Example 4.3 has characteristic polynomial $(t - 1)^2$. Its powers satisfy $S^n e_2 = n e_1 + e_2$, so they are unbounded in any norm on this finite-dimensional space. An invariant positive norm would keep $\|S^n e_2\|$ constant. This supplies a second proof of the obstruction and explains why a unit-circle eigenvalue condition does not control the representation.

A.5 Dropping continuity at a dense core

Define $q : \mathbb{R} \rightarrow \mathbb{R}$ by $q(x) = 0$ for rational x and $q(x) = -1$ otherwise. It is nonnegative on a dense subset and negative elsewhere. This function is not a quadratic form. It refutes the general continuity-free inference for arbitrary real-valued functions, not a theorem specifically asserting continuity of a finite-dimensional quadratic form. In applications to unbounded forms, the relevant continuity has to be proved in the form topology.

A.6 Dropping the gap scale

Example 7.1 has a positive spectral gap for every finite n , but no uniform lower bound. A claim that each approximation has a gap is weaker than a claim that the residual divided by that gap tends to zero. The distinction persists in infinite-dimensional problems and is especially relevant when a ground-state eigenvalue is numerically extremely small.

A.7 Dropping holomorphic control

For $n \geq 1$, set $f_n(z) = 1 + z^2/n$ on the complex plane. The functions tend locally uniformly to the nonzero function 1, while their zeros escape to infinity. This illustrates that local uniform convergence does not preserve the total number of zeros in the whole plane. Hurwitz applies on compact regions and transfers the absence of zeros there. A spectral determinant argument needs that local statement together with the specified arithmetic limit, rather than an informal matching of complete zero lists.

B An explicit comparison checklist

A proposed bridge in APP is recorded by answering the following mathematical questions. The list is a specification of a theorem, not an approval procedure.

1. What are the scalar field, test space, topology, and involution?
2. What are the source object and the map from tests into it?
3. Which form is pulled back, and with what sign and normalization?
4. Is the source sign theorem proved independently on a specified domain?
5. Which kernel or radical is removed, and is that the intended arithmetic quotient?
6. For operators, what is the domain and which adjoint or invariance identity holds there?
7. For limits, are all forms evaluated on common tests, and what scalar defect tends to zero?
8. For a complex zero limit, what is the holomorphic domain, normalization, and nonzero limiting function?
9. Which arithmetic identity identifies the construction with the completed zeta distribution or its zero divisor?
10. Which fields are implemented, which are classical inputs, and which remain conjectural?

A map can satisfy some of these conditions without satisfying all of them. Recording that partial result is useful if its scope is kept exact. A finite-dimensional form comparison, for example, may establish an arithmetic identity on a chosen span while leaving its extension to all smooth tests open.

C Logical strength of an unrestricted residual

A smaller analytic specification need not be a logically weaker proposition. The distinction is especially easy to miss when the objects to be approximated have not been fixed. The following elementary result is included in `APP/Residual.lean` as an explicit circularity check.

Let X be any set and $Q : X \rightarrow \mathbb{R}$ any function. For model values $m_n : X \rightarrow \mathbb{R}$, a lower comparison residual consists of functions $\epsilon_n : X \rightarrow \mathbb{R}$ such that, for each fixed x , $\epsilon_n(x) \rightarrow 0$ and eventually

$$m_n(x) - \epsilon_n(x) \leq Q(x).$$

Suppose also that $m_n(x) \geq 0$ for every n, x . A limit argument gives $Q(x) \geq 0$, just as in Proposition 6.1.

Proposition C.1. *The function Q is nonnegative if and only if there exists some family of nonnegative model values with a lower comparison residual as above.*

Proof. The reverse implication is the sign-transfer argument. For the forward implication, assume $Q \geq 0$, set $m_n = Q$ for every n , and set $\epsilon_n = 0$. The comparison is then equality and the error converges to zero. $\square$

The forward witness is not a finite arithmetic approximation. The proposition deliberately imposes no finite-dimensionality or arithmetic definition on the models. It shows why an unrestricted existential approximation certificate can conceal the original target. One obtains a genuine research task only after prescribing the model class and the maps, so that the constant target family is not silently used as the construction.

For example, a prescribed finite correspondence algebra might give m_n through an ordinary trace of a square. That establishes model nonnegativity by independent algebra. The remaining

comparison is then a statement about a concrete arithmetic error, whose summands and cutoff parameters can be studied. Its difficulty may still be RH-strength. The progress is the exposure of the error and its proof obligations, not an assertion that the ultimate proposition has become easier merely by changing notation.

The same caution applies to a proposed universal purity structure. If its definition includes the assertion that every zeta zero is represented in a positive space satisfying the weight identity, existence already implies RH. This can be a useful sufficient geometric conjecture when its objects have an independent construction. It cannot be presented as an unconditional imported theorem merely because all consequences of the definition are routine.

This distinction guides the stopping condition for a research iteration. A failed proof should leave a named mathematical object, a precisely typed comparison, and an estimate or construction statement that can be attacked on its own. It should not leave a new opaque proposition whose only explanation is that it would imply the old one. APP's residual declaration records the actual target function, actual model values, one-sided bound, and limiting error separately. The number-theoretic instance is still missing and is stated as missing.

References

- [1] M. Long and the YonedaAI Collaboration, *Arithmetic positivity structures*, Arithmetic Positivity Program, Part I, 2026. Companion manuscript.
- [2] M. Long and the YonedaAI Collaboration, *Finite Weil approximations*, Arithmetic Positivity Program, Part II, 2026. Companion manuscript.
- [3] A. Connes and C. Consani, Weil positivity and trace formula, the archimedean place, *Selecta Mathematica* 27 (2021), article 77. [arXiv:2006.13771](#).
- [4] N. Aronszajn, Theory of reproducing kernels, *Transactions of the American Mathematical Society* 68 (1950), 337-404. [doi:10.1090/S0002-9947-1950-0051437-7](#).
- [5] M. Suzuki, *Weil's quadratic form via the screw function*, 2026. [arXiv:2606.09096](#).
- [6] A. Connes, C. Consani, and H. Moscovici, *Zeta spectral triples*, 2025. [arXiv:2511.22755](#).
- [7] A. Connes, *The Riemann hypothesis: Past, present and a letter through time*, 2026. [arXiv:2602.04022](#).
- [8] A. Connes and W. D. van Suijlekom, Quadratic forms, real zeros and echoes of the spectral action, *Communications in Mathematical Physics* 406 (2025), article 312. [arXiv:2511.23257](#).